EMC²
where information lives®

EMC Infrastructure for Virtual Desktops

Enabled by EMC VNX Series, VMware vSphere 4.1, VMware View 4.5, and VMware View Composer 2.5

Proven Solution Guide

EMC Unified Storage Solutions

EMC² PROVEN SOLUTIONS

vmware®

Copyright © 2011 EMC Corporation. All rights reserved.

Published March, 2011

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is”. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

VMware, ESX, vMotion, VMware vCenter, VMware View, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part number: h8147.1

Table of Contents

Chapter 1: About this Document 4
Overview 4
Audience and purpose 5
Scope 5
Reference architecture 6
Hardware and software resources 8
Prerequisites and supporting documentation 10
Terminology 11
Chapter 2: VMware View Infrastructure 12
Overview 12
VMware View 4.5 13
vSphere 4.1 infrastructure 15
Windows infrastructure 16
Chapter 3: Storage Design 17
Overview 17
EMC VNX series storage architecture 18
Chapter 4: Network Design 22
Overview 22
Considerations 23
VNX for file network configuration 24
ESX network configuration 25
Cisco 6509 configuration 26
Fibre Channel network configuration 27
Chapter 5: Installation and Configuration 28
Overview 28
VMware components 29
Storage components 32
Chapter 6: Testing and Validation 35
Overview 35
Testing overview 36
Validated environment profile 36
Result analysis 38
Boot storm results 38
View refresh results 43
View recompose results 51
Antivirus results 57
Patch install results 62
Login VSI results 68

Chapter 1: About this Document

Overview

Introduction

EMC's commitment to consistently maintain and improve quality is led by the Total Customer Experience (TCE) program, which is driven by Six Sigma methodologies. As a result, EMC has built Customer Integration Labs in its Global Solutions Centers to reflect realworld deployments in which TCE use cases are developed and executed. These use cases provide EMC with an insight into the challenges that are currently facing its customers.

This Proven Solution Guide summarizes a series of best practices that were discovered, validated, or otherwise encountered during the validation of the EMC Infrastructure for Virtual Desktops Enabled by EMC® VNX™ Series, VMware vSphere™ 4.1, VMware View™ 4.5, and VMware® View Composer 2.5 solution by using the following products:

- EMC VNX series
- VMware View Manager 4.5
- VMware View Composer 2.5
- VMware vSphere 4.1
- EMC PowerPath® Virtual Edition

Use case definition

The following six use cases are examined in this solution:

- Boot storm
- View refresh operation
- View recompose operation
- Antivirus scan
- Security patch install
- User workload simulated with Login VSI tool

Chapter 6: Testing and Validation contains the test definitions and results for each use case.

Contents

This chapter contains the following topics:

Topic	See Page
Audience and purpose	5
Scope	5
Reference architecture	6
Hardware and software resources	8
Prerequisites and supporting documentation	10
Terminology	11

Audience and purpose

Audience

The intended audience for this Proven Solution Guide is:

- EMC and VMware customers
- EMC and VMware partners
- Internal EMC and VMware personnel

Purpose

The purpose of this use case is to provide a virtualized solution for virtual desktops that is powered by VMware View 4.5, View Composer 2.5, VMware vSphere 4.1, EMC VNX series, EMC VNX FAST VP, VNX FAST Cache, and storage pools.

This solution includes all the attributes required to run this environment, such as hardware and software, including Active Directory, and the required VMware View configuration.

Information in this document can be used as the basis for a solution build, white paper, best practices document, or training.

Scope

Scope

This Proven Solution Guide contains the results of testing the EMC Infrastructure for Virtual Desktops Enabled by EMC VNX Series, VMware vSphere 4.1, VMware View 4.5, and VMware View Composer solution. The objectives of this testing are to establish:

- A reference architecture of validated hardware and software that permits easy and repeatable deployment of the solution.
- The storage best practices to configure the solution in a manner that provides optimal performance, scalability, and protection in the context of the midtier enterprise market.

Not in scope

Implementation instructions are beyond the scope of this document. Information on how to install and configure VMware View 4.5 components, vSphere 4.1, and the required EMC products is out of the scope of this document. Links to supporting documentation for these products are supplied where applicable.

Reference architecture

Corresponding reference architecture

This Proven Solution Guide has a corresponding Reference Architecture document that is available on EMC Powerlink® and EMC.com. The *EMC Infrastructure for Virtual Desktops Enabled by EMC VNX Series, VMware vSphere 4.1, VMware View 4.5, and VMware View Composer 2.5 — Reference Architecture* provides more details.

If you do not have access to the document, contact your EMC representative.

The reference architecture and the results in this Proven Solution Guide are valid for 500 Windows 7 virtual desktops conforming to the workload described in Observed workload on page 36.

Reference architecture logical diagram

The following diagram depicts the logical architecture of the midsize solution.

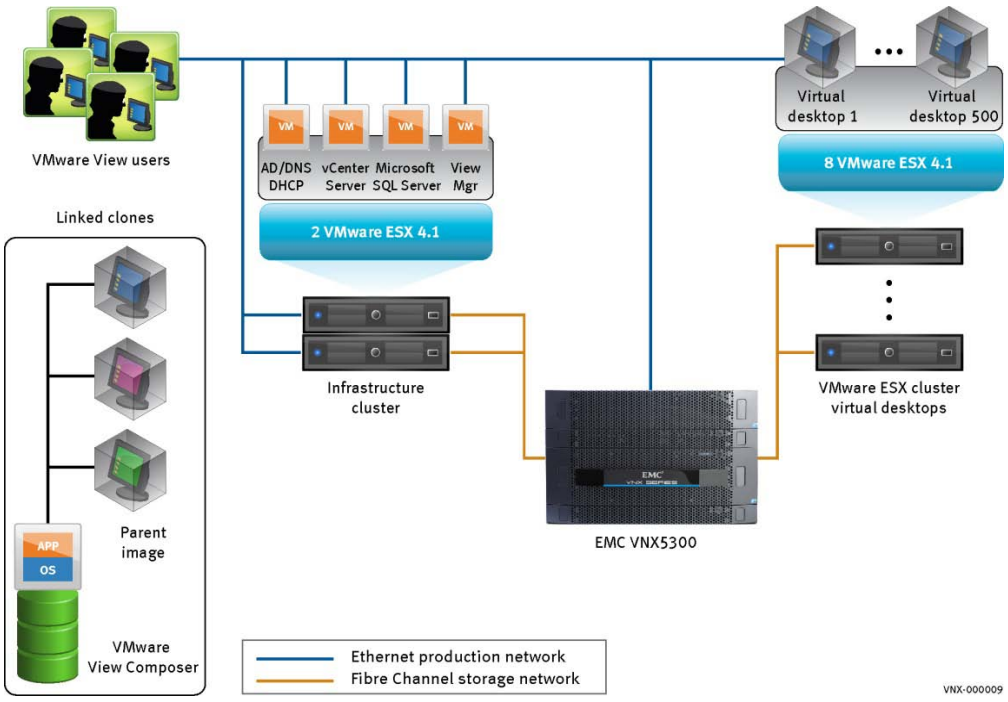

Introduction to the new EMC VNX series for unified storage

The EMC VNX series is a collection of new unified storage platforms that unifies EMC Celerra® and EMC CLARiiON® into a single product family. This innovative series meets the needs of environments that require simplicity, efficiency, and performance while keeping up with the demands of data growth, pervasive virtualization, and budget pressures. Customers can benefit from the new VNX features such as:

- Next-generation unified storage, optimized for virtualized applications
- Automated tiering with Flash and Fully Automated Storage Tiering for Virtual Pools (FAST VP) that can be optimized for the highest system performance and lowest storage cost simultaneously
- Multiprotocol support for file, block, and object with object access through Atmos™ Virtual Edition (Atmos VE)
- Simplified management with EMC Unisphere™ for a single management framework for all NAS, SAN, and replication needs
- Up to three times improvement in performance with the latest Intel multicore CPUs, optimized for Flash
- 6 Gb/s SAS back end with the latest drive technologies supported:
  - 3.5" 100 GB and 200 GB Flash, 3.5" 300 GB, and 600 GB 15k or 10k rpm SAS, and 3.5" 2 TB 7.2k rpm NL-SAS
  - 2.5" 300 GB and 600 GB 10k rpm SAS
- Expanded EMC UltraFlex™ I/O connectivity—Fibre Channel (FC), Internet Small Computer System Interface (iSCSI), Common Internet File System (CIFS), Network File System (NFS) including parallel NFS (pNFS), Multi-Path File System (MPFS), and Fibre Channel over Ethernet (FCoE) connectivity for converged networking over Ethernet

The VNX series includes five new software suites and two new software packages, making it easier and simpler to attain the maximum overall benefits.

Software suites available

- VNX FAST Suite—Automatically optimizes for the highest system performance and the lowest storage cost simultaneously.
- VNX Local Protection Suite—Practices safe data protection and repurposing.
- VNX Remote Protection Suite—Protects data against localized failures, outages, and disasters.
- VNX Application Protection Suite—Automates application copies and proves compliance.
- VNX Security and Compliance Suite—Keeps data safe from changes, deletions, and malicious activity.

Software packages available

- VNX Total Protection Pack—Includes local, remote, and application protection suites.
- VNX Total Efficiency Package—Includes all five software suites (not available for the VNX5100).
- VNX Total Value Package—Includes all three protection software suites and the Security and Compliance Suite (the VNX5100 exclusively supports this package).

Hardware and software resources

Test results

Chapter 6: Testing and Validation provides more information on the performance results.

Hardware resources

The following table lists the hardware used to validate the solution.

Hardware	Quantity	Configuration	Notes
EMC VNX5300	1	Three DAEs configured with: • Twenty one 300 GB 15k rpm SAS disks • Fifteen 2 TB 7.2k rpm near-line SAS disks • Five 100 GB EFDs	VNX shared storage
Dell PowerEdge R710	8	• Memory: 64 GB of RAM • CPU: Dual Xeon X5550, 2.67 GHz • NIC: Quad-port Broadcom BCM5709 1000BASE-T	Virtual desktop ESX® cluster
Dell PowerEdge 2950	2	• Memory: 16 GB of RAM • CPU: Dual Xeon 5160, 3 GHz • NIC: Gigabit quad-port Intel VT	Infrastructure virtual machines (VMware vCenter™ Server, DNS, DHCP, Active Directory, and Routing and Remote Access Service (RRAS)
Cisco 6509	1	• WS-6509-E switch • WS-x6748 1 Gb line cards • WS-SUP720-3B supervisor	Host connections distributed over two line cards
Brocade DS5100	2	Twenty four 8 Gb ports	Redundant SAN A or SAN B configuration
QLogic HBA	1	• Dual-port QLE2462 • Port 0 – SAN A • Port 1 – SAN B	One dual-port HBA per server connected to both the fabrics
Desktop/ virtual machines	Each	• Windows 7 Enterprise 32-bit • Memory: 768 MB • CPU: 1 vCPU • NIC: e1000 (connectivity)	

Software resources

The following table lists the software used to validate the solution.

Software	Configuration
EMC VNX5300	
VNX OE for file	Release 7.0
VNX OE for block	Release 31
ESX servers	
ESX	ESX 4.1
vCenter Server	
OS	Windows 2008 R2
VMware vCenter Server	4.1
VMware View Manager	4.5
VMware View Composer	2.5
PowerPath Virtual Edition	5.4 SP2
Desktops/virtual machines **Note**: These softwares are used to generate the test load.	
OS	MS Windows 7 Enterprise (32-bit)
VMware tools	8.3.2
Microsoft Office	Office 2007 SP2
Internet Explorer	8.0.7600.16385
Adobe Reader	9.1.0
McAfee Virus Scan	8.7.0i Enterprise

Prerequisites and supporting documentation

Technology

It is assumed that the reader has a general knowledge of the following products:

- VMware vSphere 4.1
- View Composer 2.5
- VMware View 4.5
- EMC VNX series

Supporting documents

The following documents, located on Powerlink, provide additional, relevant information. Access to these documents is based on your login credentials. If you do not have access to the following content, contact your EMC representative.

- *EMC Infrastructure for Virtual Desktops Enabled by EMC VNX Series, VMware vSphere 4.1, VMware View 4.5, and VMware View Composer 2.5—Reference Architecture*
- *Deploying Microsoft Windows 7 Virtual Desktops with VMware View—Applied Best Practices Guide*
- *EMC Performance Optimization for Microsoft Windows XP for the Virtual Desktop Infrastructure—Applied Best Practices*

VMware documents

The following documents are available on the VMware website:

- *Introduction to VMware View Manager*
- *VMware View Manager Administrator Guide*
- *VMware View Architecture Planning Guide*
- *VMware View Installation Guide*
- *VMware View Integration Guide*
- *VMware View Reference Architecture*
- *Storage Deployment Guide for VMware View*
- *VMware View Windows XP Deployment Guide*
- *VMware View Guide to Profile Virtualization*

Terminology

Introduction This section defines the terms used in this document.

Term	Definition
EMC VNX FAST Cache	EMC VNX FAST Cache is a feature that enables the use of EFD as an expanded cache layer for the array.
FAST VP	FAST VP is a pool-based feature of the VNX series that supports scheduled migration of data to different storage tiers based on the performance requirements of individual 1 GB slices in a storage pool.
Linked clone	A virtual desktop created by VMware View Composer from a writeable snapshot paired with a read-only replica of a master image.
Login VSI	A third-party benchmarking tool developed by Login Consultants that simulates real-world VDI workload by using an AutoIT script and determines the maximum system capacity based on the response time of the users.
Replica	A read-only copy of a master image used to deploy linked clones.
VMware View Composer	Integrates effectively with VMware View Manager to provide advanced image management and storage optimization.

Chapter 2: VMware View Infrastructure

Overview

Introduction

The general design and layout instructions described in this chapter apply to the specific components used during the development of this solution.

Contents

This chapter contains the following topics:

Topic	See Page
VMware View 4.5	13
vSphere 4.1 infrastructure	15
Windows infrastructure	16

VMware View 4.5

Introduction

VMware View delivers rich and personalized virtual desktops as a managed service from a virtualization platform built to deliver the entire desktop, including the operating system, applications, and user data. With VMware View 4.5, desktop administrators can virtualize the operating system, applications, and user data, and deliver modern desktops to end users. VMware View 4.5 provides centralized automated management of these components with increased control and cost savings. VMware View 4.5 improves business agility while providing a flexible high-performance desktop experience for end users across a variety of network conditions.

Deploying VMware View components

This solution is deployed using a single View Manager server instance that is capable of scaling up to 2,000 virtual desktops. Deployments of up to 10,000 virtual desktops are possible by using multiple View Manager servers.

The core elements of a VMware View 4.5 implementation are:

- View Manager Connection Server
- View Composer 2.5
- vSphere 4.1

Additionally, the following components are required to provide the infrastructure for a View 4.5 deployment:

- Microsoft Active Directory
- Microsoft SQL Server
- DNS Server
- DHCP Server

View Manager Connection Server

The View Manager Connection Server is the central management location for virtual desktops and has the following key roles:

- Broker connections between users and virtual desktops
- Control the creation and retirement of virtual desktop images
- Assign users to desktops
- Control the state of the virtual desktops
- Control access to the virtual desktops

View Composer 2.5

View Composer 2.5 works directly with vCenter Server to deploy, customize, and maintain the state of the virtual desktops when using linked clones. The tiered storage capabilities of View Composer 2.5 allow the read-only replica and the linked clone disk images to be on the dedicated storage. This allows superior scaling in large configurations.

View Composer linked clones

VMware View with View Composer 2.5 uses the concept of linked clones to quickly provision virtual desktops. This solution uses the tiered storage feature of View Composer to build linked clones and place their replica images on separate data stores as shown in the following diagram:

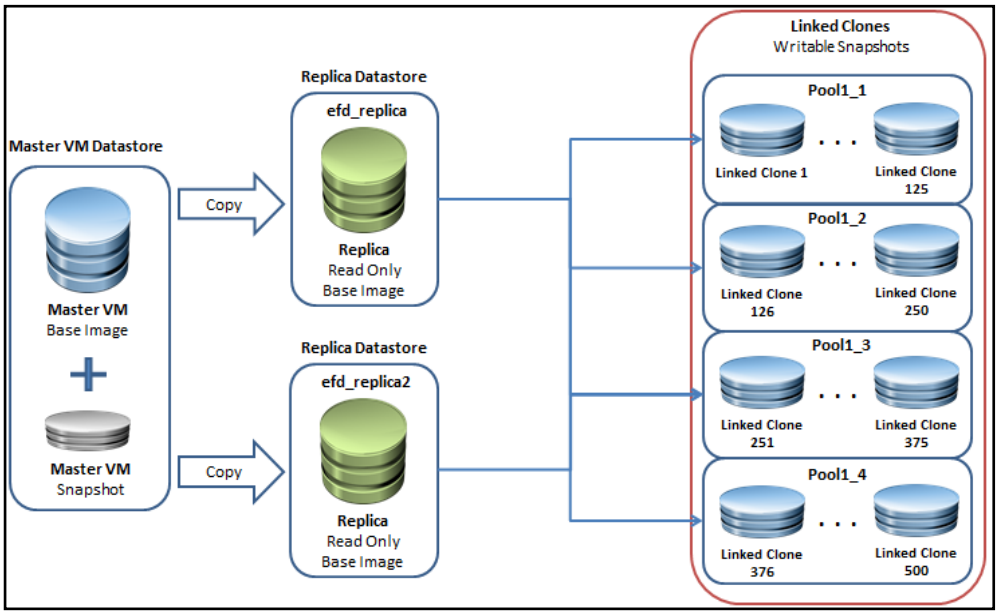

In this configuration, the operating system reads from the common read-only replica image and writes to the linked clone. Any unique data created by the virtual desktop is also stored in the linked clone. A logical representation of this relationship is shown in the following diagram:

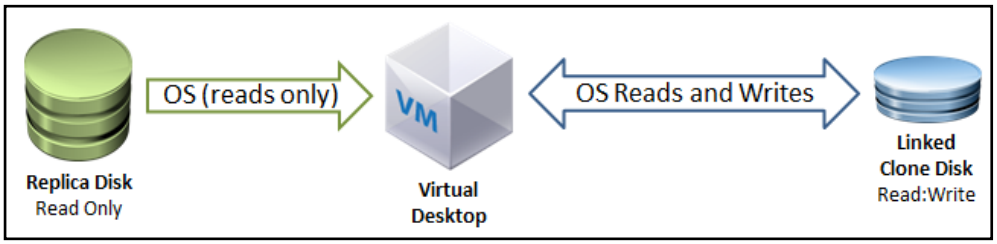

Automated pool configuration

In this solution, all 500 desktops were deployed on two automated desktop pools by using a common Windows 7 master image. Dedicated data stores were used for the replica images and the linked clone storage. The linked clones were distributed across the following four data stores:

- efd_replica and efd_replica2 contained the read-only copies of the master Windows 7 image. These data stores were backed by two 100 GB EFDs. Each replica image supported 250 linked clones.
- Pool1_1 through Pool1_4 were used to store the linked clones. Each desktop pool was configured to use all four Pool1_x data stores so that all the virtual desktops were evenly distributed across the data stores.

vSphere 4.1 infrastructure

vSphere 4.1 overview

VMware vSphere 4.1 is the market-leading virtualization hypervisor used across thousands of IT environments around the world. VMware vSphere 4.1 can transform or virtualize computer hardware resources, including CPU, RAM, hard disk, and network controller, to create a fully functional virtual machine that runs its own operating system and applications just like a physical computer.

The high-availability features in VMware vSphere 4.1 along with VMware Distributed Resource Scheduler (DRS) and Storage vMotion® enable seamless migration of virtual desktops from one ESX server to another with minimal or no impact to customers' usage.

vCenter Server cluster

The following diagram shows the cluster configuration from vCenter Server.

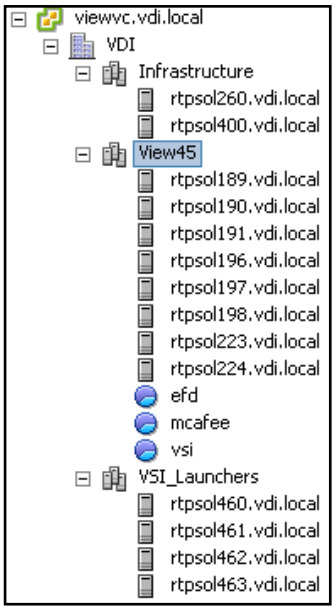

The Infrastructure cluster holds the following virtual machines:

- Windows 2008 R2 domain controller—provides DNS, Active Directory, and DHCP services.
- Windows 2008 R2 SQL 2008—provides databases for vCenter Server, View Composer, and other services in the environment.
- Windows 2003 R2 View 4.5—provides services for managing virtual desktops.
- Windows 2008 R2 vCenter Server—provides management services for the VMware clusters and View Composer.
- Windows 7 Key Management Service (KMS)—provides a method to activate Windows 7 desktops.

The View 4.5 cluster (called View45 in the diagram) consists of 500 virtual desktops.

Windows infrastructure

Introduction

Microsoft Windows provides the infrastructure used to support the virtual desktops and includes the following components:

- Microsoft Active Directory
- Microsoft SQL Server
- DNS Server
- DHCP Server

Microsoft Active Directory

The Windows domain controller runs the Active Directory service that provides the framework to manage and support the virtual desktop environment. Active Directory has several functions such as the following:

- Manage the identities of users and their information
- Apply group policy objects
- Deploy software and updates

Microsoft SQL Server

Microsoft SQL Server is a relational database management system (RDBMS). SQL Server 2008 is used to provide the required databases to vCenter Server and View Composer.

DNS Server

DNS is the backbone of Active Directory and provides the primary name resolution mechanism of Windows servers and clients.

In this solution, the DNS role is enabled on the domain controller.

DHCP Server

The DHCP Server provides the IP address, DNS Server name, gateway address, and other information to the virtual desktops.

In this solution, the DHCP role is enabled on the domain controller.

Chapter 3: Storage Design

Overview

Introduction

The storage design described in this chapter applies to the specific components of this solution.

Contents

This chapter contains the following topic:

Topic	See Page
EMC VNX series storage architecture	18

EMC VNX series storage architecture

Introduction

The EMC VNX series is a dedicated network server optimized for file and block access that delivers high-end features in a scalable and easy-to-use package.

The VNX series delivers a single-box block and file solution, which offers a centralized point of management for distributed environments. This makes it possible to dynamically grow, share, and cost-effectively manage multiprotocol file systems and provide multiprotocol block access. Administrators can take advantage of the simultaneous support for NFS and CIFS protocols by allowing Windows and Linux/UNIX clients to share files by using the sophisticated file-locking mechanisms of VNX for file and VNX for block for high-bandwidth or for latency-sensitive applications.

This solution uses both block-based and file-based storage to leverage the benefits that each of the following provides:

- Block-based storage over Fibre Channel (FC) is used to store the VMDK files for all virtual desktops. This has the following benefits:
  - Block storage leverages the VAAI APIs (introduced in vSphere 4.1) that includes a hardware-accelerated copy to improve the performance and for granular locking of the VMFS to increase scaling.
  - PowerPath Virtual Edition (PowerPath/VE) allows better performance and scalability as compared to the native multipathing options.
- File-based storage is provided by a CIFS export. This has the following benefits:
  - Redirection of user data and roaming profiles to a central location for easy backup and administration.
  - Single instancing and compression of unstructured user data to provide the highest storage utilization and efficiency.

This section explains the configuration of the storage that was provided over FC to the ESX cluster to store the VMDK images and the storage that was provided over CIFS to redirect user data and roaming profiles

Storage layout

The following diagram shows the storage layout of the disks.

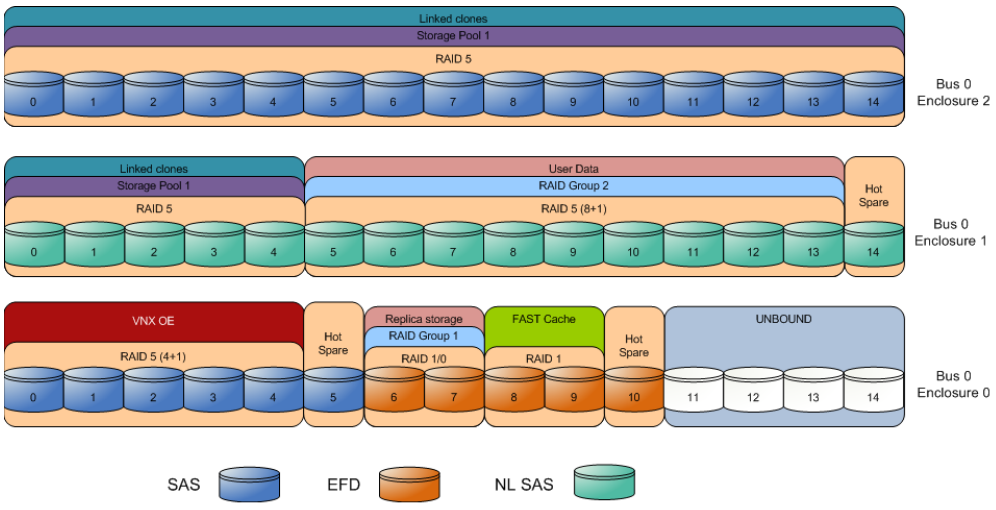

Storage layout overview

The following storage configuration was used in the solution:

- SAS disks (0_0 to 0_4) are used for the VNX OE.
- Disks 0_5, 0_10, and 1_14 are hot spares.
- EFDs (0_6_ and 0_7) on the RAID 1/0 group are used to store the linked clone replicas.
- EFDs (0_8 and 0_9) are used for EMC VNX FAST Cache. There are no user-configurable LUNs on these drives.
- SAS disks (2_0 to 2_14) and near-line SAS disks (1_0 to 1_4) on the RAID 5 pool are used to store linked clones. The storage pool uses FAST with SAS and near-line SAS disks to optimize both performance and capacity across the pool. FAST Cache is enabled for the entire pool. Four LUNs of 750 GB each are created from the pool and presented to the ESX servers.
- Near-line SAS disks (1_5 to 1_13) on the RAID 5 (8+1) group are used to store user data and roaming profiles. Two file systems are created on two LUNs, one for profiles and the other for user data.
- SAS disks (0_11 to 0_14) are unbound. They are not used for validation tests.

Please note that this reference architecture has been developed using RAID 5 in order to maximize performance. Customers, specifically those using 1 TB or larger drives, whose goal is maximum availability during drive rebuilds (should that occur in their environment) should choose RAID 6, because of the benefit of the additional parity drive.

EMC VNX FAST Cache

VNX FAST Cache, a part of the VNX FAST suite, enables EFDs to be used as an expanded cache layer for the array. The VNX5300 is configured with two 100 GB EFDs in a RAID 1 configuration for a 93 GB read/write capable cache. This is the minimum amount of FAST Cache. Larger configurations are supported for scaling beyond 500 desktops.

FAST Cache has array-wide features available for both file and block storage. FAST Cache works by examining 64 KB chunks of data in FAST Cache enabled objects on the array. Frequently accessed data is copied to the FAST Cache and subsequent accesses to that data chunk are serviced by FAST Cache. This allows immediate promotion of very active data to the EFDs. This dramatically improves the response time for very active data and reduces the data hot spots that can occur within the LUN.

FAST Cache is an extended read/write cache that can absorb read-heavy activities such as boot storms and antivirus scans, and write-heavy workloads such as operating system patches and application updates.

EMC VNX FAST VP

FAST is a pool-based feature available for VNX for block LUNs that migrates data to different storage tiers based on the performance requirements of the data.

The pool1_x LUNs are built on a storage pool configured on RAID 5 with a mix of SAS and near-line SAS drives. Initially, the linked clones are placed on the SAS tier. The data created by the linked clones that is not frequently accessed is automatically migrated to the near-line SAS storage tier. This releases space in the faster SAS tier for more active data.

EMC PowerPath Virtual Edition

Each data store that is used to store VMDK files is placed on the VNX5300 storage over FC. PowerPath/VE is enabled for all FC-based LUNs to efficiently use all the available paths for storage and to minimize the effect of micro-bursting I/O patterns.

vCenter Server storage layout

The data store configuration in vCenter Server is as follows:

- pool 1_1 through pool 1_4—Each of the 750 GB data stores accommodates 125 users. This allows each desktop to grow to a maximum average size of 6 GB. The pool of desktops created in View Manager is balanced across all these data stores.
- efd_replica and efd_replica2—These data stores are on two 100 GB EFDs with RAID 1/0. The input/output to these LUNs is strictly read-only except during operations that require copying a new replica into the data store.

VNX shared file systems

Virtual desktops use two VNX shared file systems, one for user profiles and the other to redirect user storage. Each file system is exported to the environment through a CIFS share.

The following table shows the file systems used for user profiles and redirected user storage.

File system	Use	Size
profiles_fs	Users' profile data	1 TB
userdata1_fs	Users' data	2 TB

EMC VNX for File Home Directory feature

The EMC VNX for File Home Directory feature uses the userdata1_fs file system to automatically map the H: drive of each virtual desktop to the users' own dedicated subfolder on the share. This ensures that each user has a dedicated home drive share with exclusive rights to that folder. This export does not need to be created manually. The Home Directory feature automatically maps this for each user.

The Documents folder of the users is also redirected to this share. This allows users to recover the data in the Documents folder by using the VNX Snapshots for File. The file system is set at an initial size of 1 TB. However, it can extend itself automatically when more space is required.

Profile export

The profiles_fs file system is used to store user roaming profiles. It is exported through CIFS. The UNC path to the export is configured in Active Directory for roaming profiles as shown in the following figure:

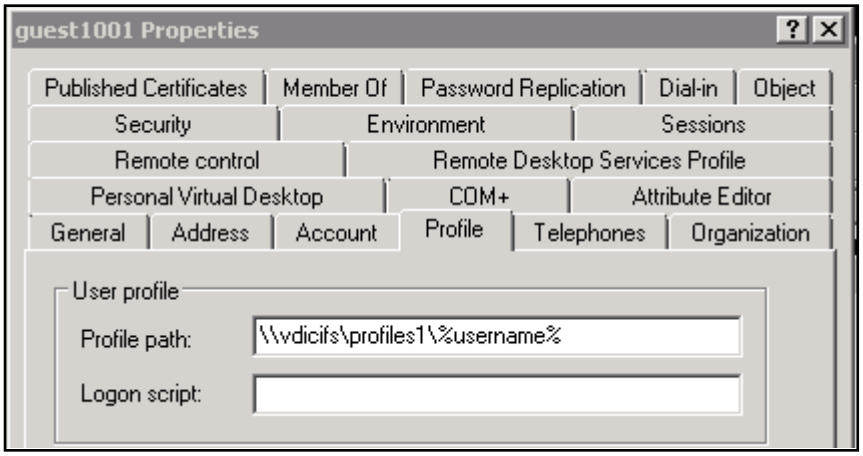

Capacity

The file systems leverage Virtual Provisioning™ and compression to provide flexibility and increased storage efficiency. If single instancing and compression are enabled, unstructured data such as user documents typically leads to a 50 percent reduction in consumed storage.

The VNX file systems for user profiles and documents are configured as follows:

- profiles_fs is configured to consume 1 TB of space. Assuming 50 percent space savings, each profile can grow up to 4 GB in size. The file system can be extended if more space is needed.
- userdata_fs is configured to consume 1 TB of space. Assuming 50 percent space savings, each user will be able to store 4 GB of data. The file system can be extended if more space is needed.

Chapter 4: Network Design

Overview

Introduction

This chapter describes the network design used in this solution.

Contents

This chapter contains the following topics:

Topic	See Page
Considerations	23
VNX for file network configuration	24
ESX network configuration	25
Cisco 6509 configuration	26
Fibre Channel network configuration	27

Considerations

Physical design considerations	EMC recommends that switches support gigabit Ethernet (GbE) connections and Link Aggregation Control Protocol (LACP), and the ports on switches support copper-based media.
Logical design considerations	This validated solution uses virtual local area networks (VLANs) to segregate network traffic of various types to improve throughput, manageability, application separation, high availability, and security. The IP scheme for the virtual desktop network must be designed such that there are enough IP addresses in one or more subnets for the DHCP Server to assign them to each virtual desktop.
Link aggregation	VNX platforms provide network high availability or redundancy by using link aggregation. This is one of the methods used to address the problem of link or switch failure. Link aggregation is a high-availability feature that enables multiple active Ethernet connections to appear as a single link with a single MAC address and potentially multiple IP addresses. In this solution, LACP is configured on VNX, which combines two GbE ports into a single virtual device. If a link is lost in the Ethernet port, the link fails over to another port. All network traffic is distributed across the active links.

VNX for file network configuration

Data Mover ports

VNX5300 consists of two Data Movers. These Data Movers can be configured in an active/active or active/passive configuration. In this solution, the Data Movers operated in the active/passive mode. In the active/passive configuration, the passive Data Mover serves as a failover device for the active Data Mover.

The VNX5300 Data Mover was configured for four 1 Gb interfaces on a single SLIC. Link Aggregation Control Protocol (LACP) was used to configure ports cge-2-0 and cge-2-1. Ports cge-2-2 and cge-2-3 were left free for further expansion.

The lacp1 device was used to support virtual machine traffic, home folder access, and external access for roaming profiles.

The external_interface device was used for administrative purposes to move data in and out of the private network on VLAN 274. Both the interfaces exist on the LACP1 device configured on cge-2-0 and cge-2-1

The ports are configured as follows:

```
external_interface protocol=IP device=lacp1
        inet=10.6.121.55 netmask=255.255.255.0
        broadcast=10.6.121.255
        UP, Ethernet, mtu=1500, vlan=521,
        macaddr=0:60:48:1b:76:92
lacp1_int protocol=IP device=lacp1
        inet=192.168.80.5 netmask=255.255.240.0
        broadcast=192.168.95.255
        UP, Ethernet, mtu=9000, vlan=274,
        macaddr=0:60:48:1b:76:92
```

LACP configuration on the Data Mover

To configure the link aggregation that uses cge-2-0 and cge-2-1 on server_2, type the following at the command prompt:

```
$ server_sysconfig server_2 -virtual -name <Device Name> -
create trk –option "device=cge-2-0,cge-2-1 protocol=lacp"
```

To verify if the ports are channeled correctly, type the following:

```
$ server_sysconfig server_2 -virtual -info lacp1
server_2:
*** Trunk lacp1: Link is Up ***
*** Trunk lacp1: Timeout is Short ***
*** Trunk lacp1: Statistical Load C is IP ***
Device   Local Grp  Remote Grp    Link   LACP Duplex Speed
-----------------------------------------------------------------
cge-2-0     10003       5888       Up    Up   Full 1000 Mbs
cge-2-1     10003       5888       Up    Up   Full 1000 Mbs
```

The remote group number must match for both the ports and the LACP status must be “Up.” Verify if the appropriate speed and duplex are established as expected.

ESX network configuration

ESX NIC teaming

All network interfaces in this solution use 1 GbE connections. The Dell R710 servers use four on-board Broadcom GbE Controllers for all network connections. The following diagram shows the vSwitch configuration in vCenter Server.

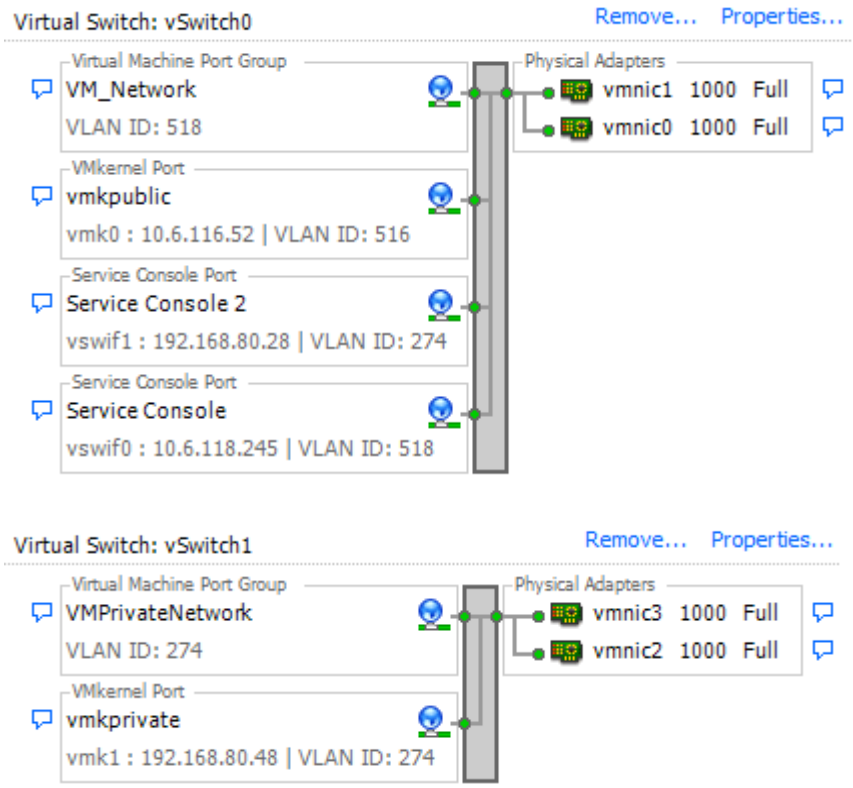

The following table lists the configured port groups in vSwitch0 and vSwitch1.

Virtual switch	Configured port groups	Used to
vSwitch0	VM_Network	Provide external access for administrative virtual machines.
vSwitch0	Vmkpublic	Mount NFS data stores on the public network for OS installation and patch installs.
vSwitch0	Service Console 2	Manage private network administration traffic.
vSwitch0	Service Console	Manage public network administration traffic.
vSwitch1	VMPrivateNetwork	Provide a network connection for virtual desktops and LAN traffic.
vSwitch1	Vmkprivate	Mount multiprotocol exports from the VNX system on the private VLAN for administrative purposes.

Cisco 6509 configuration

Overview

The nine-slot Cisco Catalyst 6509-E switch provides high port densities that are ideal for many wiring closet, distribution, and core network deployments as well as data center deployments.

Cabling

In this solution, the ESX server and VNX Data Mover cabling are evenly spread across two WS-x6748 1 Gb line cards to provide redundancy and load balancing of the network traffic.

Server uplinks

The server uplinks to the switch are configured in a port channel group to increase the utilization of server network resources and to provide redundancy. The vSwitches are configured to load balance the network traffic on the originating port ID.

The following is an example of the configuration for one of the server ports:

```
description 8/10 9048-43 rtpsol189-1
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 274,516-527
 switchport mode trunk
 no ip address
 spanning-tree portfast trunk
```

Data Movers

The network ports for each VNX5300 Data Mover are connected to the 6509-E switch. The server_2 ports cge-2-0 and cge-2-1 are configured with LACP, which provides redundancy in case of a NIC or port failure.

The following is an example of the switch configuration for one of the Data Mover ports:

```
description 7/4 9047-4 rtpsol22-dm2.0
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 274,516-527
 switchport mode trunk
 mtu 9216
 no ip address
 spanning-tree portfast trunk
 channel-group 23 mode active
```

Fibre Channel network configuration

Introduction

Two Brocade DS5100 series FC switches are used to provide the storage network for this solution. The switches are configured in a SAN A/SAN B configuration to provide a fully redundant fabric.

Each server has a single connection to each fabric to provide load-balancing and failover capabilities. Each storage processor has two links to the SAN fabrics for a total of four available front-end ports. The zoning is configured so that each server has four available paths to the storage array.

Zone configuration

Single initiator and multiple target zoning are used in this solution. Each server initiator is zoned to two storage targets on the array. The following diagram shows the zone configuration for the SAN A fabric.

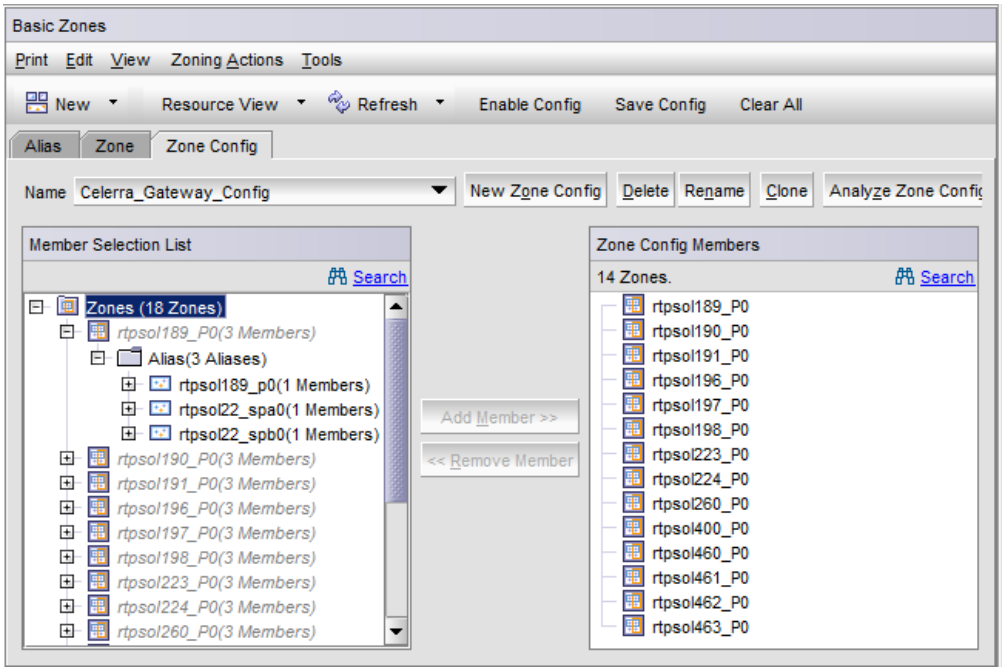

Chapter 5: Installation and Configuration

Overview

Introduction

This chapter provides an overview of the configuration of the following components:

- Desktop pools
- Storage pools
- FAST Cache
- Auto-tiering (FAST VP)
- VNX Home Directory
- PowerPath/VE

The installation and configuration steps for the following components are available on the VMware website (www.vmware.com):

- VMware View Connection Server
- VMware View Composer 2.5
- VMware ESX 4.1
- VMware vSphere 4.1

The installation and configuration of the following components are not covered:

- Microsoft System Center Configuration Manager (SCCM)
- Microsoft Active Directory, DNS, and DHCP
- vSphere and its components
- Microsoft SQL Server 2008 R2

Contents

This chapter contains the following topics:

Topic	See Page
VMware components	29
Storage components	32

VMware components

VMware View installation overview

The *VMware View Installation Guide* available on the VMware website has detailed procedures to install View Connection Server and View Composer 2.5. There are no special configuration instructions required for this solution.

The *ESX Installable and vCenter Server Setup Guide* available on the VMware website has detailed procedures to install vCenter Server and ESX and is not covered in further detail in this paper. There are no special configuration instructions required for this solution.

VMware View setup

Before deploying the desktop pools, ensure that the following steps from the *VMware View Installation Guide* have been completed:

- Prepare Active Directory
- Install View Composer 2.5 on vCenter Server
- Install View Connection Server
- Add a vCenter Server instance to View Manager

VMware View desktop pool configuration

VMware recommends using a maximum of 250 desktops per replica image, which requires creating a unique pool for every 250 desktops. In this solution, persistent automated desktop pools were used.

To create a persistent automated desktop pool as configured for this solution, complete the following steps:

1. Log in to the **VMware View Administration** page, which is located at https://server/admin, where "server" is the IP address or DNS name of the View Manager server.
2. Click the **Pools** link in the left pane.
3. Click **Add** under the **Pools** banner.
4. In the **Type** page, select **Automated Pool** and click **Next**.
5. In the **User Assignment** page, select **Dedicated** and ensure that the **Enable automatic assignment** checkbox is selected. Click **Next**.
6. In the **vCenter Server** page, select **View Composer linked clones** and select a vCenter Server that supports View Composer, as shown in the following figure. Click **Next**.

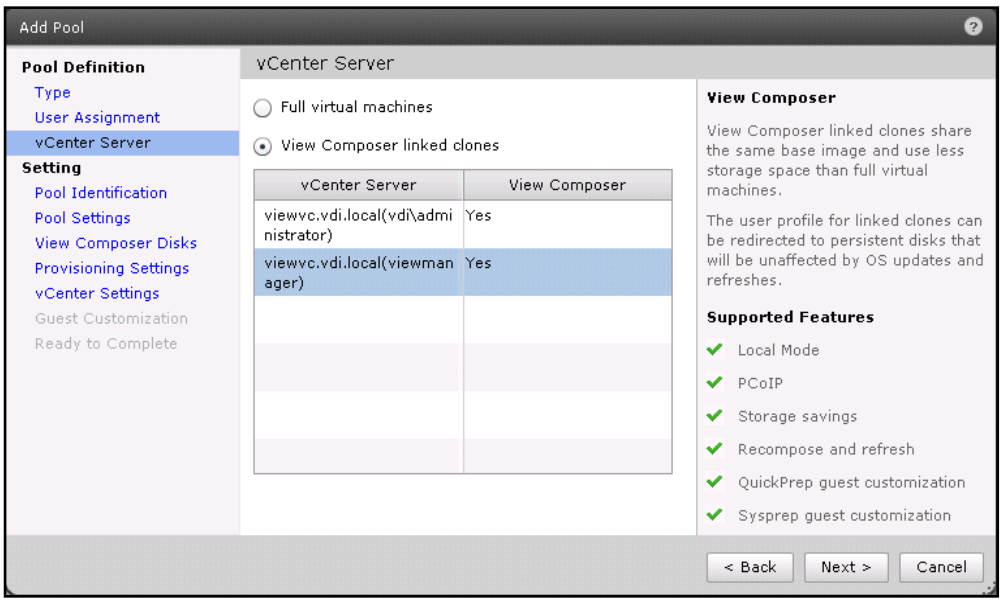

7. In the **Pool Identification** page, enter the required information and click **Next**.
8. In the **Pool Settings** page, make any required changes and click **Next**.
9. In the **View Composer Disks** page, select **Do not redirect Windows profile** and click **Next**.
10. In the **Provisioning Settings** page, select a name for the desktop pool and enter the number of desktops to provision, as shown in the following figure. Click **Next**.

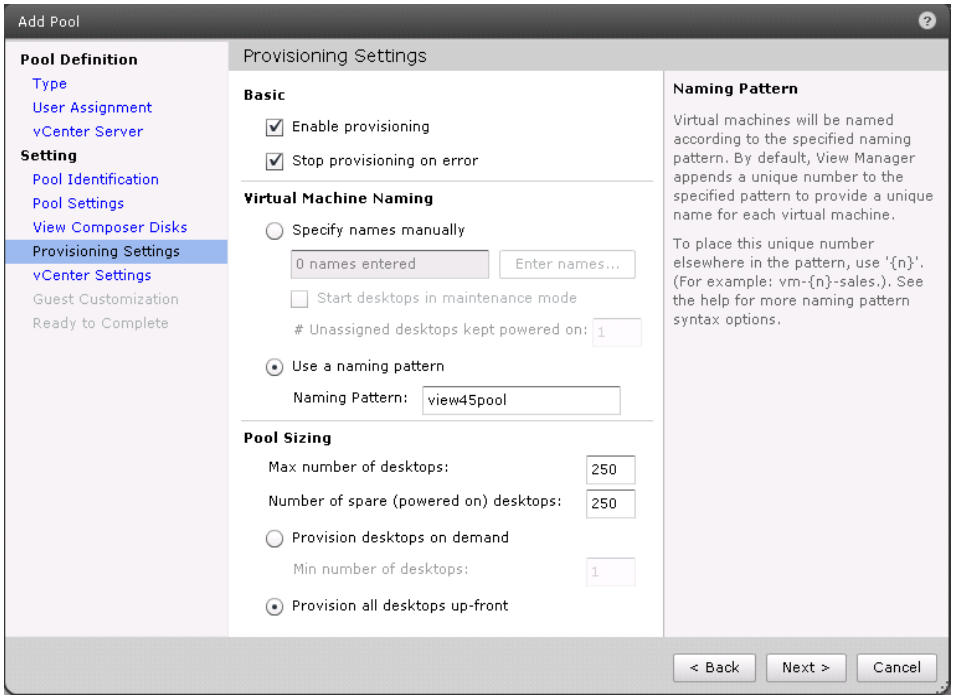

11. In the **vCenter Settings** page, browse to select a default image, a folder for the virtual machines, the cluster hosting the virtual desktops, the resource

pool to hold the desktops, and the data stores that will be used to deploy the desktops, and then click **Next**.

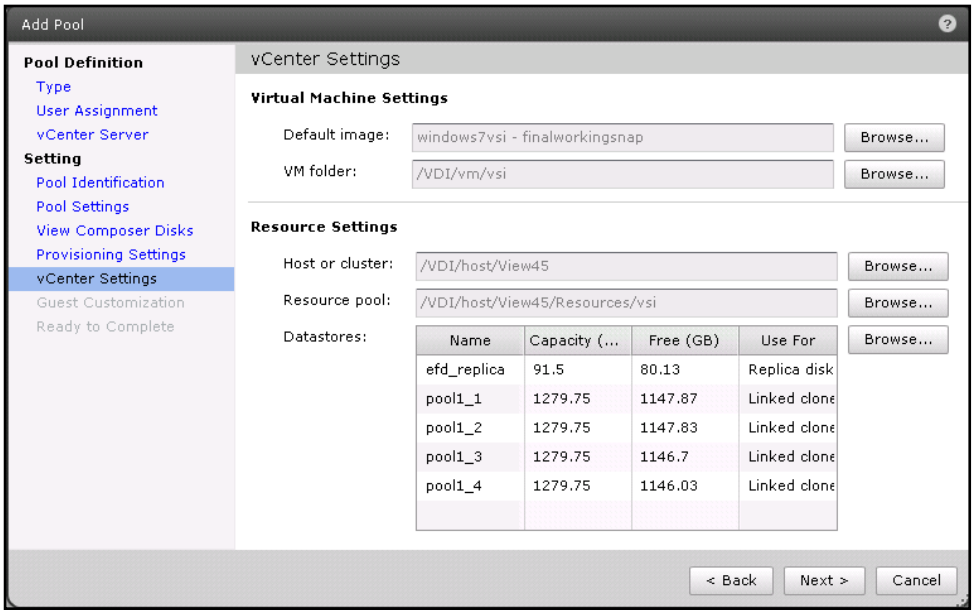

12. In the **Select Datastores** page, select **Use different datastore for View Composer replica disks** and select the data stores for replica and linked clone images, and then click **OK**.

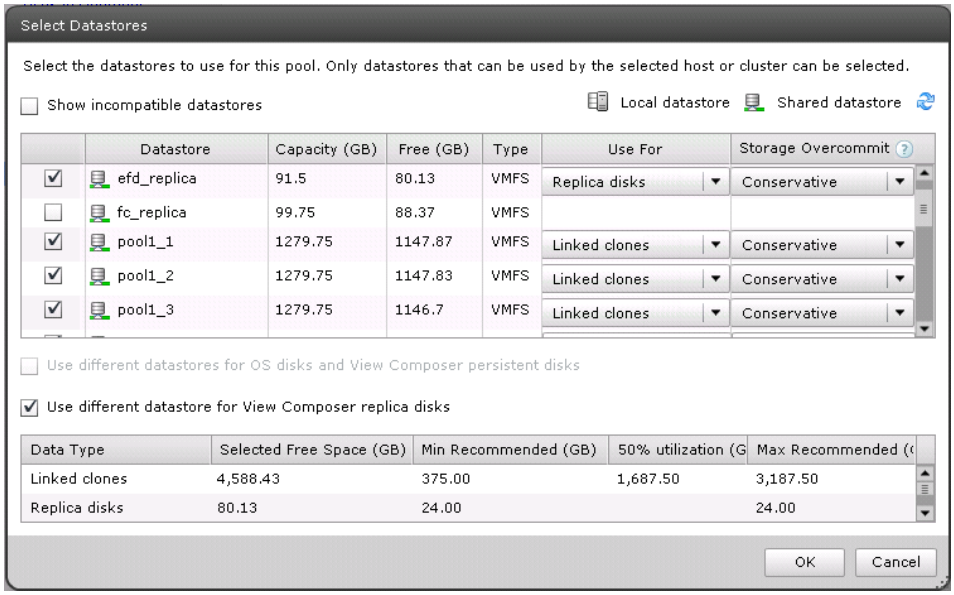

13. In the **Guest Customization** page, select the **domain** and **AD container**, and then select **Use QuickPrep**. Click **Next.**
14. In the **Ready to Complete** page, verify the settings for the pool, and then click **Finish** to start the deployment of the virtual desktops.

PowerPath Virtual Edition

PowerPath/VE 5.4.1 supports ESX 4.1. The *EMC PowerPath/VE for VMware vSphere Installation and Administration Guide* available on Powerlink provides the procedure to install and configure PowerPath/VE. There are no special configuration instructions required for this solution.

The PowerPath/VE binaries and support documentation are available on Powerlink.

Storage components

Storage pools

Storage pools in the EMC VNX OE support heterogeneous drive pools. In this solution, a 20-disk pool with RAID 5 was configured from 15 SAS disks and five near-line SAS drives. Four thick LUNs, each 750 GB in size, were created from this storage pool, as shown in the following figure. FAST Cache was enabled for the pool.

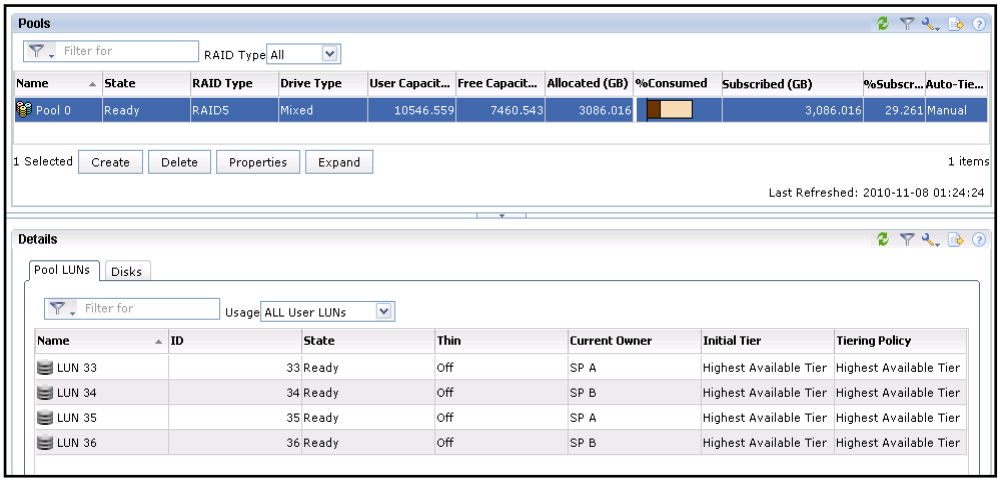
Pools
Filter for
RAID Type All
Name State RAID Type Drive Type User Capacit... Free Capacit... Allocated (GB) %Consumed Subscribed (GB) %Subscr... Auto-Tie...
Pool 0 Ready RAID5 Mixed 10546.559 7460.543 3086.016 3,086.016 29.261 Manual
1 Selected Create Delete Properties Expand 1 items
Last Refreshed: 2010-11-08 01:24:24
Details
Pool LUNs Disks
Filter for
Usage ALL User LUNs
Name ID State Thin Current Owner Initial Tier Tiering Policy
LUN 33 33 Ready Off SP A Highest Available Tier Highest Available Tier
LUN 34 34 Ready Off SP B Highest Available Tier Highest Available Tier
LUN 35 35 Ready Off SP A Highest Available Tier Highest Available Tier
LUN 36 36 Ready Off SP B Highest Available Tier Highest Available Tier

For each LUN in the storage pool, the tiering policy is set to **Highest Available Tier** to ensure that all frequently accessed desktop data remains on the SAS disks. As data ages and is used infrequently, it is moved to the near-line SAS drives in the pool.

Enable FAST Cache

FAST Cache is enabled as an array-wide feature in the system properties of the array in Unisphere™. Click the **FAST Cache** tab, click **Create**, and then select the eligible EFDs to create the FAST Cache. There are no user-configurable parameters for the FAST Cache.

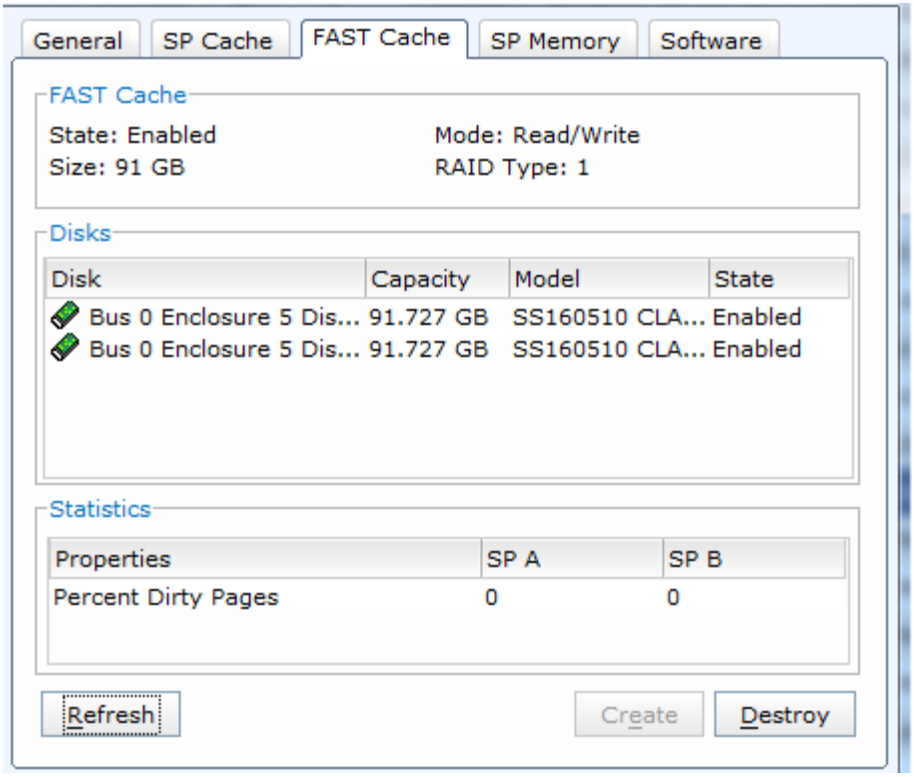

FAST Cache was not enabled for the replica storage in this solution. The replica images were serviced from the EFDs. Enabling FAST Cache for these LUNs causes additional overhead without added performance.

If the replica images are stored on SAS disks, enable FAST Cache for those LUNs. To enable FAST Cache for any LUN in a pool, go to the properties of the pool in Unisphere™ and click the **Advanced** tab. Select **Enabled** to enable FAST Cache, as shown in the following figure.

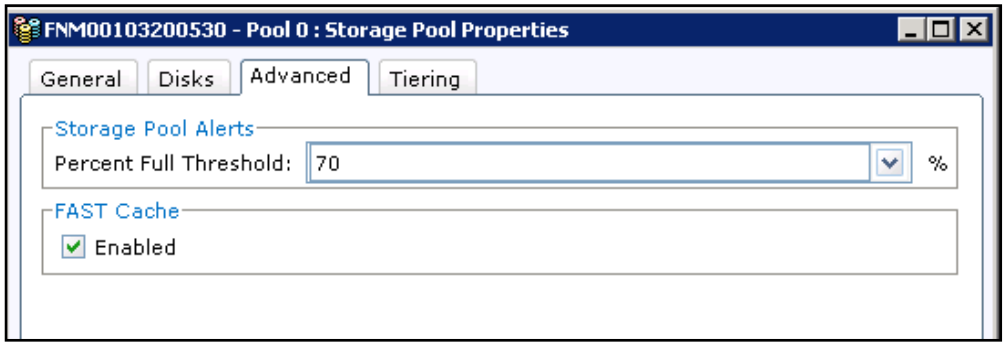

Configure FAST

To configure the FAST feature for a pool LUN, go to the properties for a pool LUN in Unisphere and click the **Tiering** tab and set the tiering policy for the LUN.

VNX Home Directory feature

The VNX Home Directory installer is available on the NAS Tools and Application CD for each VNX OE for file release and can be downloaded from Powerlink.

After the VNX Home Directory feature is installed, use the Microsoft Management Console (MMC) snap-in to configure the feature. A sample configuration is shown in the following two figures.

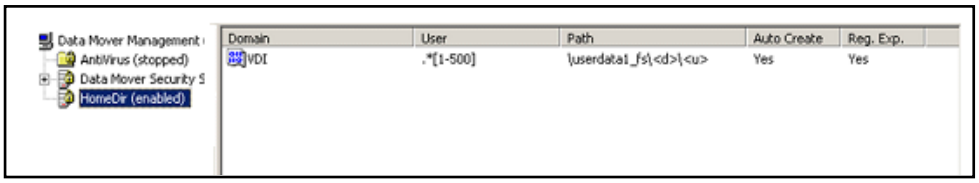

For any user account that ends with a suffix between 1 and 500, the sample configuration shown in the following figure automatically creates a user home directory in the following location and maps the H: drive to this path: \userdata1_fs file system in the format \userdata1_fs\<domain>\<user>

Each user has exclusive rights to the folder.

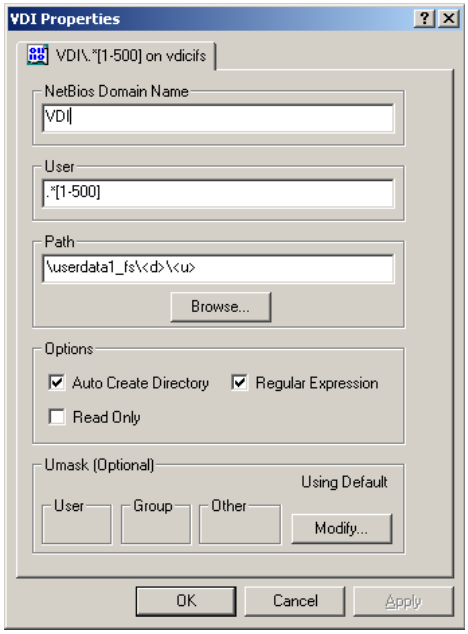

Chapter 6: Testing and Validation

Overview

Introduction

This chapter describes the tests that were run to validate the configuration of the solution.

Contents

This chapter contains the following topics:

Topic	See Page
Testing overview	36
Validated environment profile	36
Result analysis	38
Boot storm results	38
View refresh results	43
View recompose results	51
Antivirus results	57
Patch install results	62
Login VSI results	68

Testing overview

Introduction

This chapter provides a summary and characterization of the tests performed to validate the solution. The goal of the testing was to characterize the performance of the solution and component subsystems during the following scenarios:

- Boot storm of all desktops
- View desktop refresh of all desktops
- View recompose of all desktops
- McAfee antivirus full scan on all desktops
- Security patch install with Microsoft SCCM
- User workload testing using Login VSI

The steps used to configure McAfee and SCCM are beyond the scope of this document.

Validated environment profile

Observed workload

A commercial desktop workload generator was used to run an example task worker benchmark with the Windows 7 virtual desktops. The following table shows the observed workload that was used to size this reference architecture:

Windows 7 workload							
	Committed bytes	**Read IOPS**	**Write IOPS**	**Total IOPS**	**Active RAM (MB)**	**% Processor time**	**Network bytes/s**
Avg	522349163.5	3.9	5.3	9.2	264.3	7.5	75551.1
95th	589459456.0	4.0	26.4	30.4	453.0	36.6	145559.2
Max	599506944.0	577.0	875.0	1452.0	460.0	109.3	5044232.8

Traditional sizing

From the observed workload there are two traditional ways of sizing the I/O requirements, average IOPS and 95th percentile IOPS. The following table shows the number of disks required to meet the IOPS requirements by sizing for both the average and the 95th percentile IOPS:

Windows 7 disk requirements					
Avg IOPS	**# Users**	**Total IOPS**	**Read: Write Mix**	**IOPS**	**FC disks required**
9	500	4500	45:55	Read: 2000	10
				Write: 2500	13
95th IOPS	**# Users**	**Total IOPS**	**Read: Write Mix**	**IOPS**	**FC disks required**
30.4	500	15200	15:85	Read: 2280	12
				Write:12920	65

Sizing on the average IOPS can yield good performance for the virtual desktops in steady state. However, this leaves insufficient headroom in the array to absorb high peaks in I/O and the performance of the desktops will suffer during boot storms, desktop recompose or refresh tasks, antivirus DAT updates and similar events. Change management becomes the most important focus of the View administrator because all tasks must be carefully balanced across the desktops to avoid I/O storms.

To combat the issue of I/O storms, the disk I/O requirements can be sized based on the 95th percentile load. Sizing to the 95th percentile ensures that 95 percent of all values measured for IOPS fall below that value. Sizing by this method ensures great performance in all scenarios except during the most demanding of mass I/O events. However, the disadvantage of this method is cost because it takes 77 disks to satisfy the I/O requirements instead of 23 disks. This leads to higher capital and operational costs.

Use cases

Six common use cases were executed to validate whether the solution performed as expected under heavy load situations.

The following are the tested use cases tested:

- Simultaneous boot of all desktops
- View refresh operation on all desktops
- View recompose operation on all desktops
- Full antivirus scan of all desktops
- Installation of five security updates using SCCM on all desktops
- Login and steady state user load simulated using the Login VSI medium workload

In each use case, a number of key metrics are presented showing the overall performance of the solution.

Login VSI

To run a user load against the desktops, the Virtual Session Index (VSI) tool was used. VSI provided the guidance to gauge the maximum number of users a desktop environment can support. The Login VSI workload can be categorized as light, medium, heavy, and custom. A medium workload was selected for testing and had the following characteristics:

- The workload emulated a medium knowledge worker, who uses Microsoft Office, Internet Explorer, and PDF.
- After a session had started, the medium workload repeated every 12 minutes.
- The response time was measured every 2 minutes during each loop.
- The medium workload opened up to five applications simultaneously.
- The type rate was 160 ms for each character.
- The medium workload in VSI 2.0 was approximately 35 percent more resource-intensive than VSI 1.0.
- Approximately 2 minutes of idle time were included to simulate real-world users.

Each loop of the medium workload opened and used the following:

- Microsoft Outlook 2007—Browsed 10 messages.
- Internet Explorer—One instance was left open (BBC.co.uk), one instance browsed Wired.com, Lonelyplanet.com and a heavy Flash application gettheglass.com (not used with MediumNoFlash workload).

- Microsoft Word 2007—One instance to measure the response time and one instance to edit the document.
- Bullzip PDF Printer and Acrobat Reader—The Word document was printed and the PDF was reviewed.
- Microsoft Excel 2007—A very large sheet was opened and random operations were performed.
- Microsoft PowerPoint 2007—A presentation was reviewed and edited.
- 7-zip—Using the command line version, the output of the session was zipped.

Login VSI launcher

A Login VSI launcher is a Windows system that launches desktop sessions on target virtual desktops. There are two types of launchers—master and slave. There is only one master in a given test bed and there can be as many slave launchers as required.

The number of desktop sessions a launcher can run is typically limited by the CPU or memory resources. Login consultants recommend using a maximum of 45 sessions per launcher with two CPU cores (or two dedicated vCPUs) and 2 GB RAM, when the GDI limit has not been tuned (default). However with the GDI limit tuned, this limit extends to 60 sessions per two-core machine.

In this validated testing, 500 desktop sessions were launched from 12 launcher virtual machines, resulting in approximately 42 sessions established per launcher. Each launcher virtual machine is allocated two vCPUs and 4 GB of RAM. There were no bottlenecks observed on the launchers during the VSI-based tests.

FAST Cache configuration

For all tests, FAST Cache was enabled for the storage pool holding the four Pool1_x data stores. FAST Cache is not enabled for the EFD-based replica image.

Replica storage configuration

Two LUNs were created on two EFDs with a RAID 1/0 configuration for hosting replica data stores. The replicas were split across both SPs for the load balancing. Read cache is enabled on the LUNs.

Result analysis

Introduction

This section explains the results for the different test scenarios.

Boot storm results

Test methodology

This test was conducted by selecting all the desktops in vCenter Server and selecting *Power On*. Overlays are added to the graphs to show when the last power-on task completed and when the IOPS to the pool LUNs achieved a steady state. For the boot storm test, all the desktops were powered on within 8 minutes and achieved steady state approximately 3 minutes later. The total start-to-finish time was approximately 11 minutes.

EFD replica disk load

The following graph shows the IOPS from one of the EFDs that contains the replica data store.

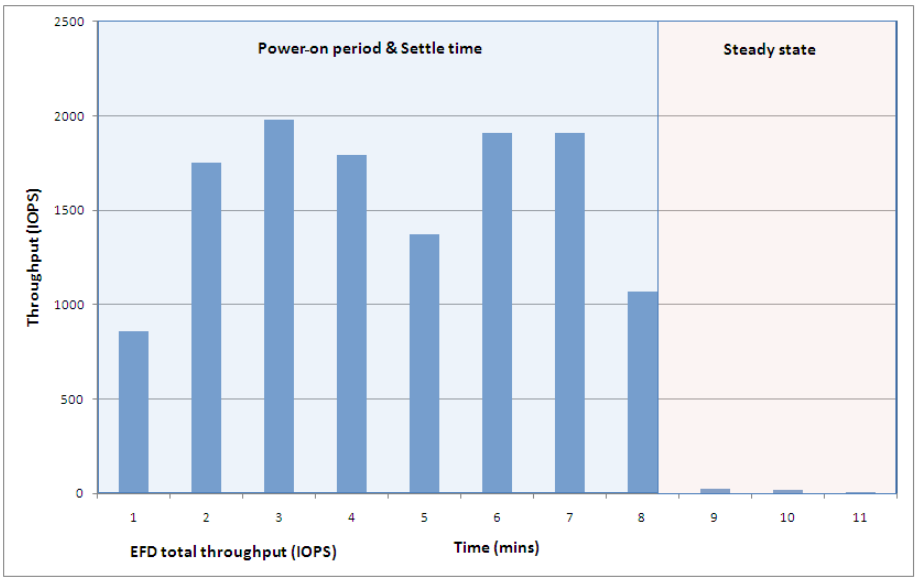

Each EFD serviced around 2,000 IOPS at peak load.

EFD replica LUN load

The following graph shows the IOPS and response time metrics from the replica LUNs.

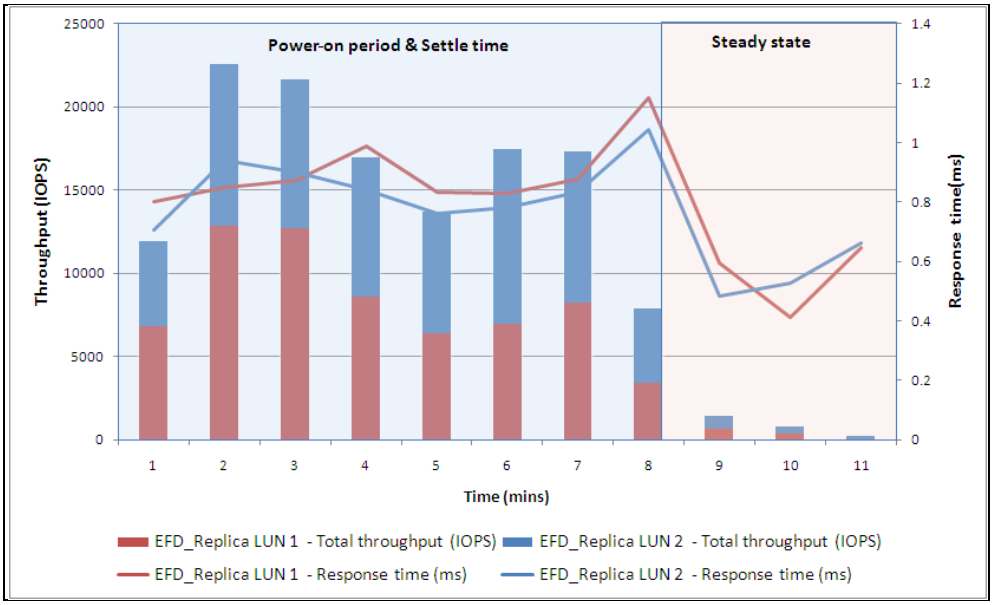

During peak load, both replica LUNs serviced a total of nearly 22,596 IOPS with a maximum response time of just over 1 ms. With read cache enabled on the EFD LUNs, the load on the EFDs was reduced during the peak I/O requirements of the boot storm.

Pool individual disk load

The following graph shows the disk IOPS for a single SAS drive in the storage pool that stores the four Pool1_x data stores. Because the statistics from all the drives in the pool were similar, a single drive is reported for the purpose of clarity and readability of the graph.

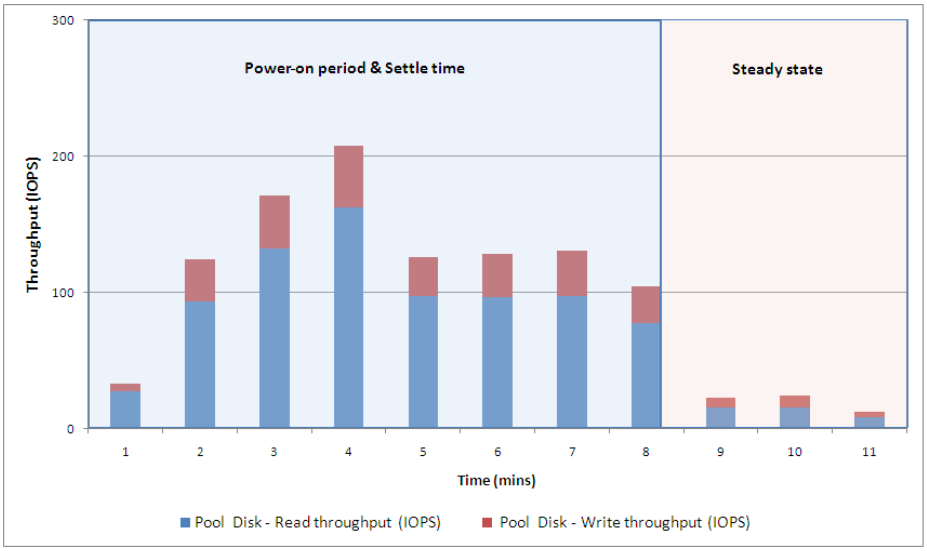

During peak load, the disk serviced a maximum of 207 IOPS. FAST Cache helped to reduce the disk load.

Pool LUN load

The following graph shows the LUN IOPS and response time from the Pool1_3 data store. Because the statistics from all pools were similar, a single pool is reported for the purpose of clarity and readability of the graph.

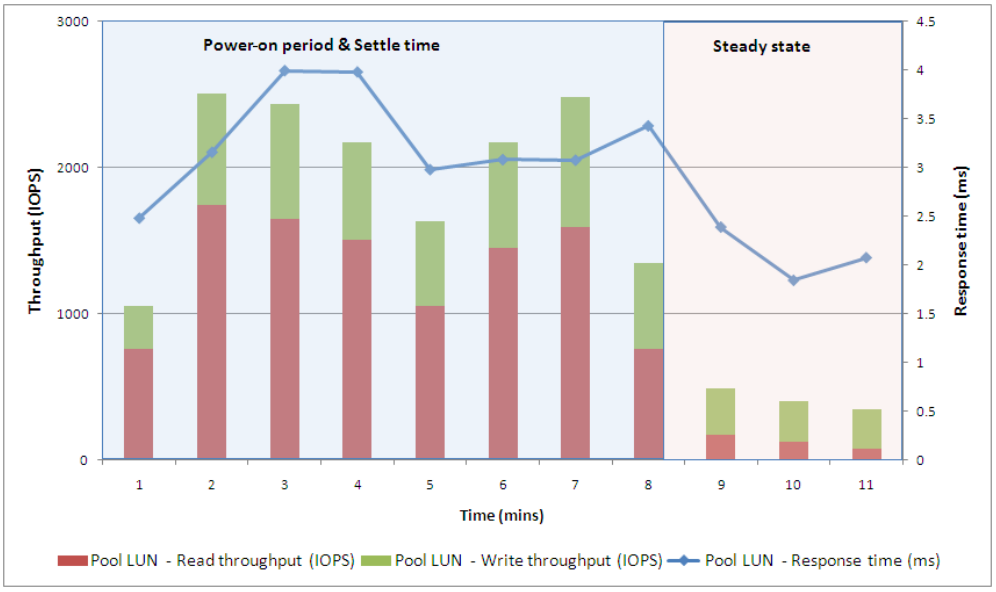

During peak load, the LUN response time did not exceed 4 ms and the data store serviced over 2,500 IOPS.

Storage processor IOPS

The following graph shows the total IOPS serviced by the storage processor during the test.

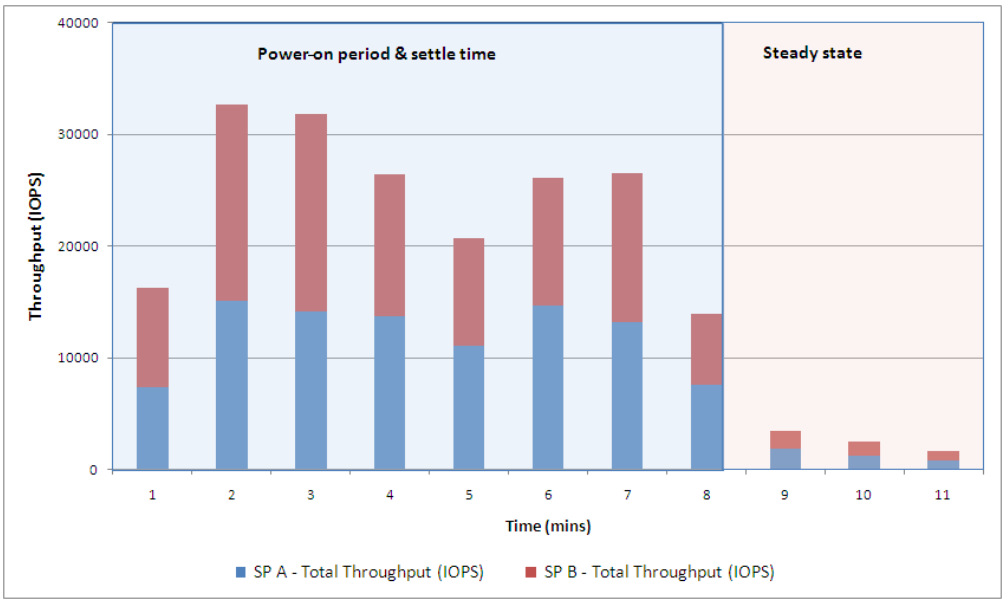

Storage processor utilization

The following graph shows the storage processor utilization during the test. The replicas were split across both SPs, which caused the load to be balanced across both SPs equally.

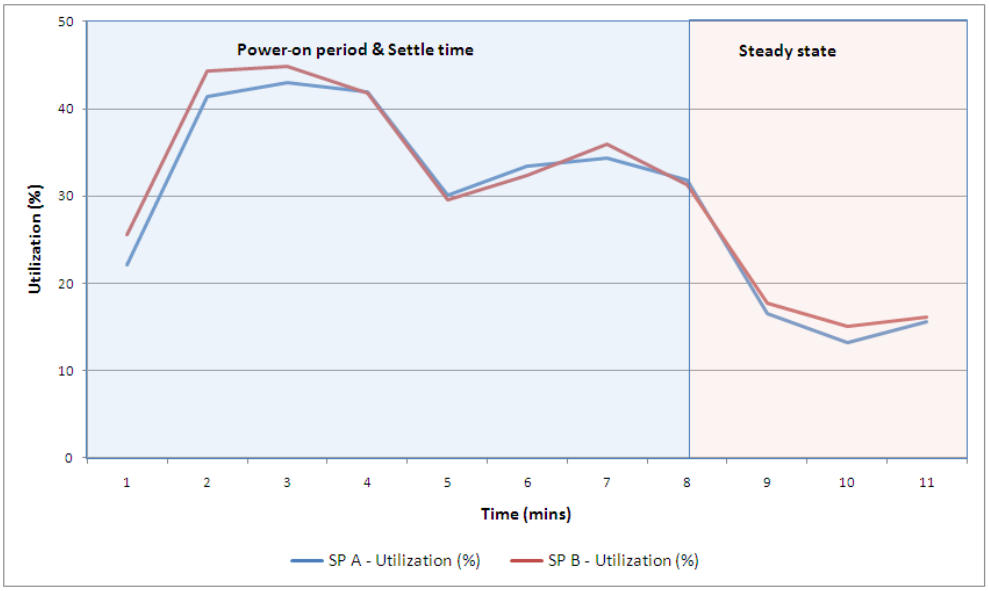

The replica traffic caused high levels of I/O during the peak load of the boot storm test while the SP utilization remained below 50 percent.

FAST Cache IOPS

The following graph shows the IOPS serviced from FAST Cache during the boot storm test.

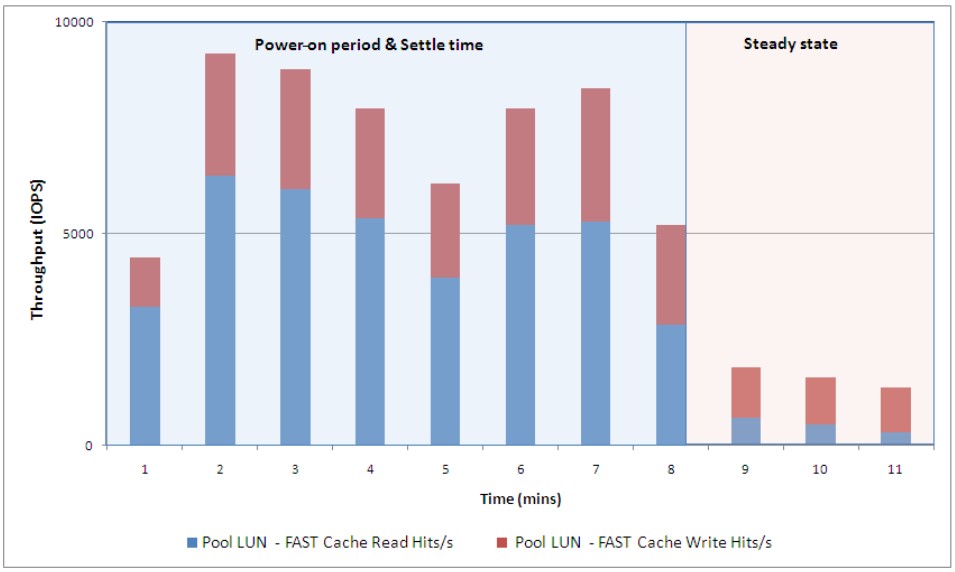

At peak load, FAST Cache serviced around 9,250 IOPS from the linked clone data stores, which is the equivalent of 47 FC disks servicing 200 IOPS each.

ESX CPU load

The following graph shows the CPU load from the ESX servers in the View cluster. All servers had similar results. Therefore, a single server is reported.

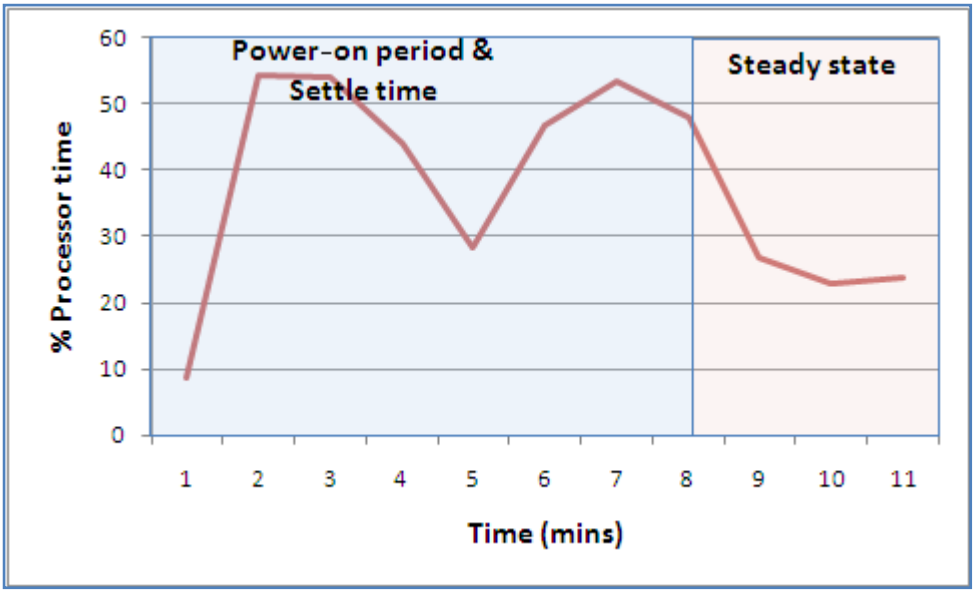

The ESX server briefly achieved a CPU utilization of approximately 55 percent during peak load in this test.

ESX disk response time

The following graph shows the Average Guest Millisecond/Command counter, which is shown as GAVG in esxtop. This counter represents the response time for I/O operations issued to the storage array. The average of both LUNs hosting the replica storage is shown in the following graph.

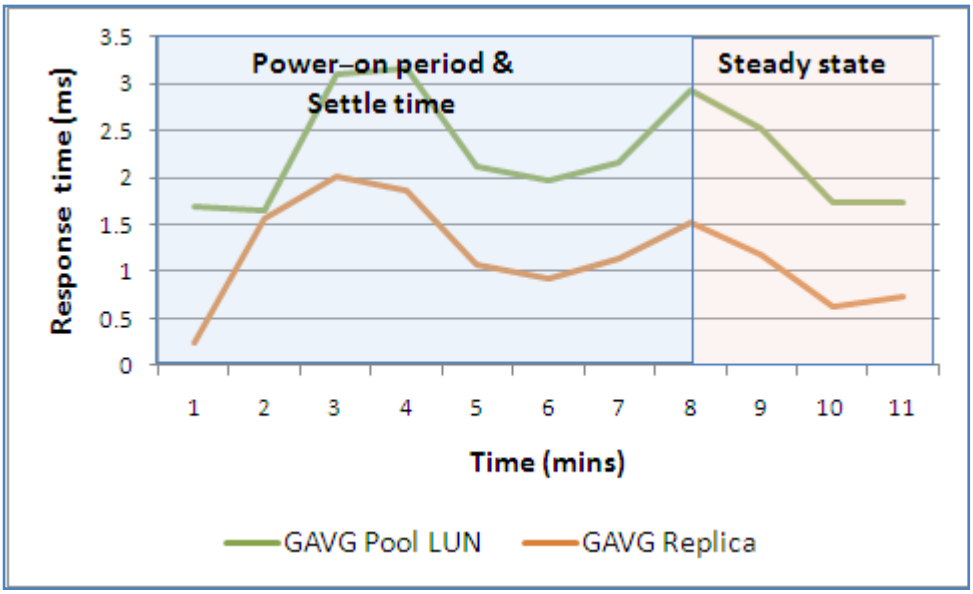

The GAVG values for the EFD replica storage and the linked clone storage on the Pool1_x data stores were below 3.5 ms. This indicates excellent performance under this load.

View refresh results

Test methodology

This test was conducted by selecting a refresh operation for all desktops in both pools from the View Manager administration console. A refresh for all desktops in one pool was started and was followed immediately by a refresh for all desktops in the other pool with only a few seconds of delay. No users were logged in during the test. Overlays are added to the graphs to show when the last power-on task completed and when the IOPS to the pool LUNs achieved a steady state.

For the refresh test, all vCenter Server tasks completed in approximately 49 minutes. The first pool completed the refresh tasks after 28 minutes and achieved steady state after 54 minutes. The second pool completed the tasks after 49 minutes and achieved steady state after 66 minutes. The start-to-finish time was approximately 66 minutes for all desktops.

EFD replica disk load

The following graph shows the IOPS from one of the EFDs that contains the replica data store.

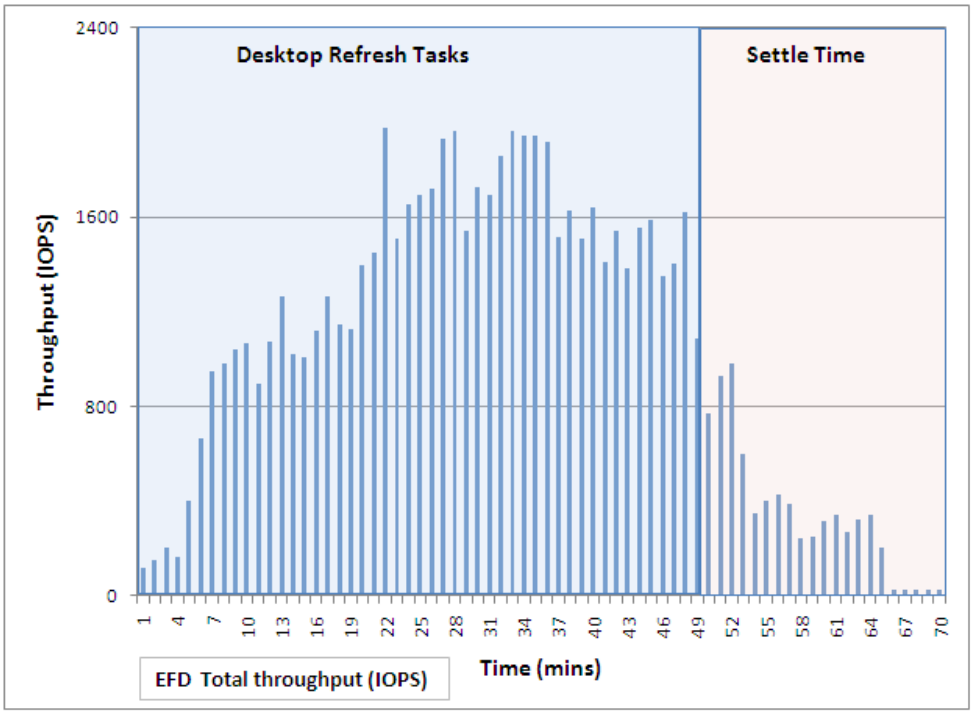

Each EFD serviced nearly 2,000 IOPS at peak load, which indicates that the disks were not driven to saturation.

EFD replica LUN load

The following graph shows the IOPS from the replica LUNs.

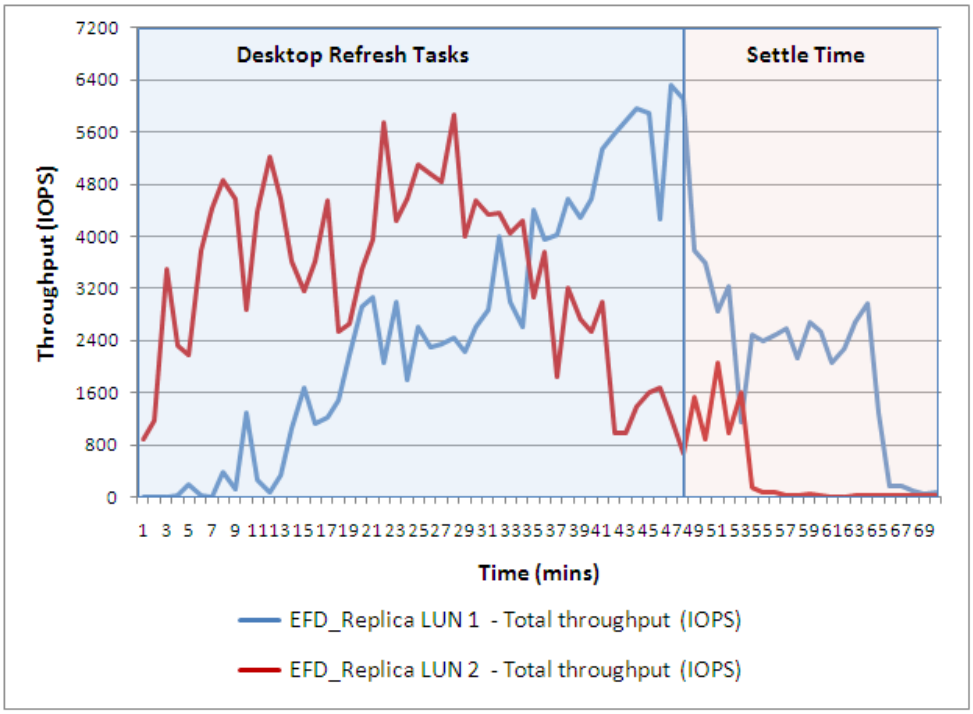

Both replica LUNs serviced a total of nearly 8,354 IOPS during peak load. The first pool achieved steady state much faster than the second pool because it completed the refresh tasks almost 20 minutes before the second pool.

EFD replica LUN response time

The following graph shows the response time of replica LUNs.

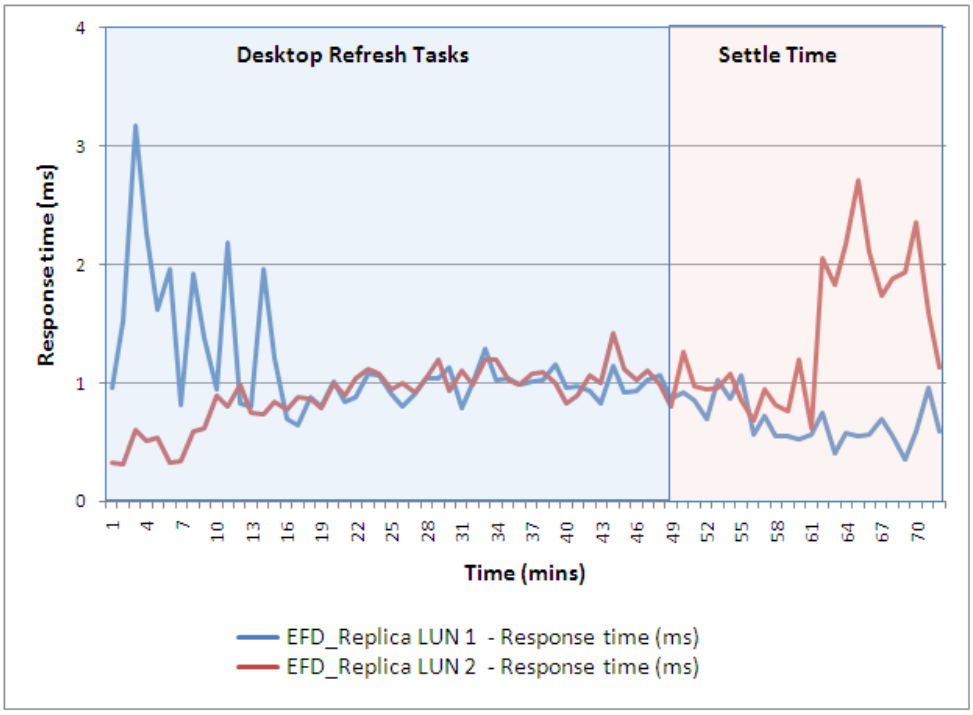
Desktop Refresh Tasks
Settle Time
Response time (ms)
Time (mins)
EFD_Replica LUN 1 - Response time (ms)
EFD_Replica LUN 2 - Response time (ms)

The peak response time values of the replica LUNs were below 3.5 ms. The first pool achieved steady state much faster than the second pool because it completed the refresh tasks almost 20 minutes before the second pool.

Pool individual disk load

The following graph shows the disk IOPS for a single SAS drive in the storage pool that stores the four Pool1_x data stores. Because the statistics for all drives in the pool were similar, only a single disk is reported.

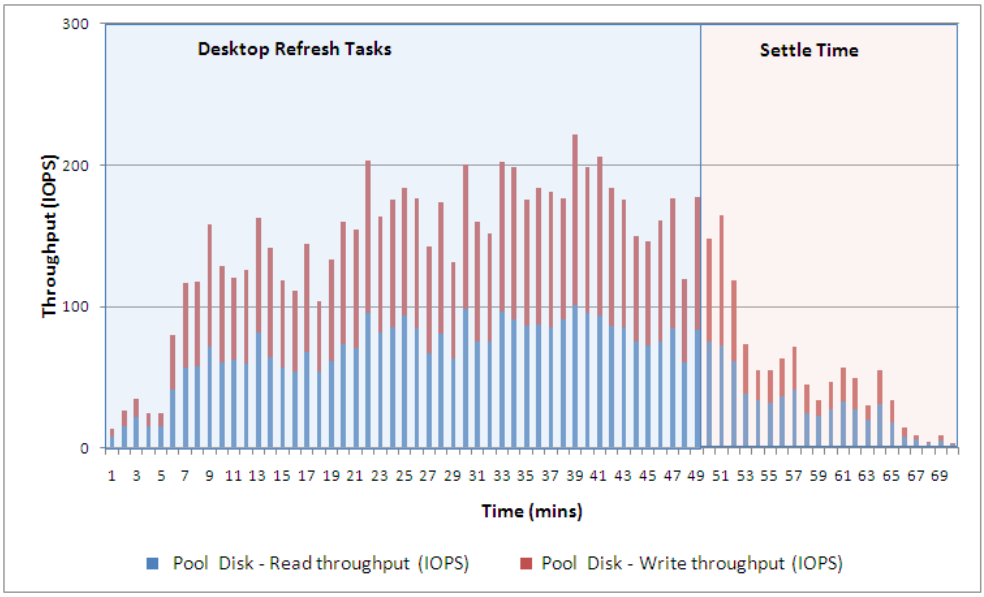

During peak load, the disk serviced a maximum of 221 IOPS. FAST Cache helped to reduce the disk load.

Pool LUN load

The following graph shows the LUN IOPS and response time from the Pool1_3 data store. Because the statistics from all pools were similar, only a single pool is reported for clarity and readability of the graph.

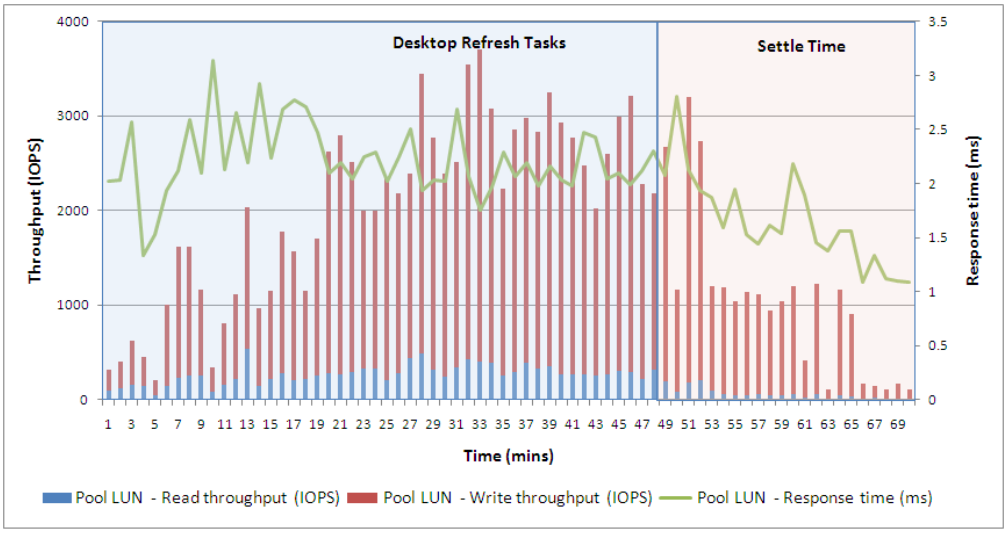

During peak load, the data store serviced over 3,700 IOPS and the LUN response time remained within 3.5 ms.

Storage processor IOPS

The following graph shows the total IOPS serviced by the storage processor during the test.

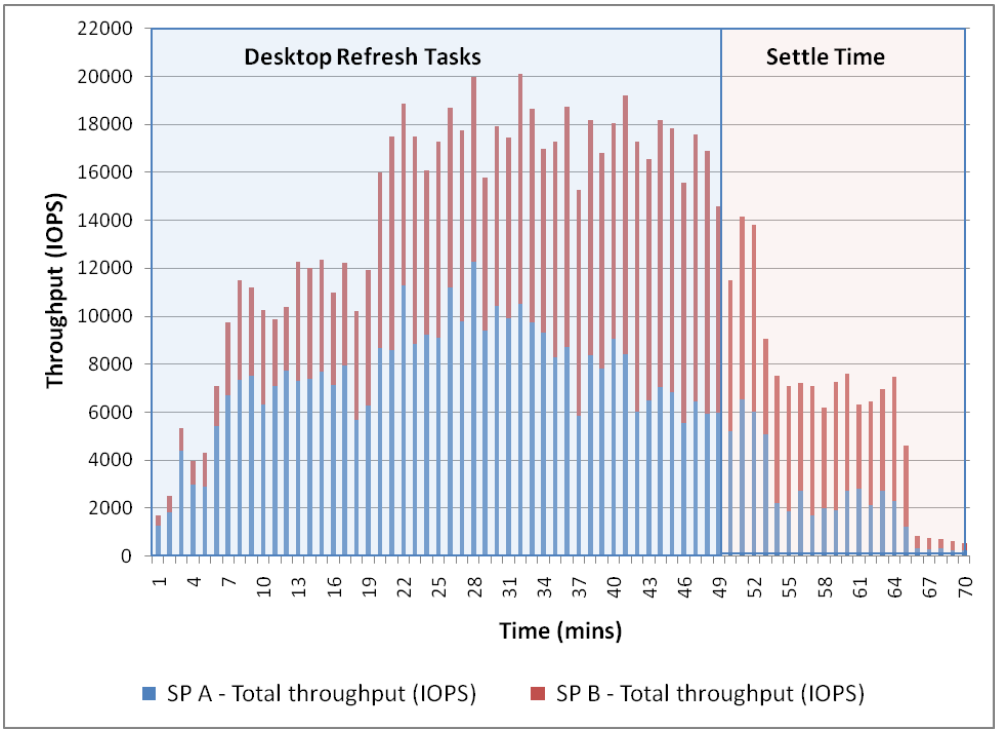

Storage processor utilization

The following graph shows the storage processor utilization during the test.

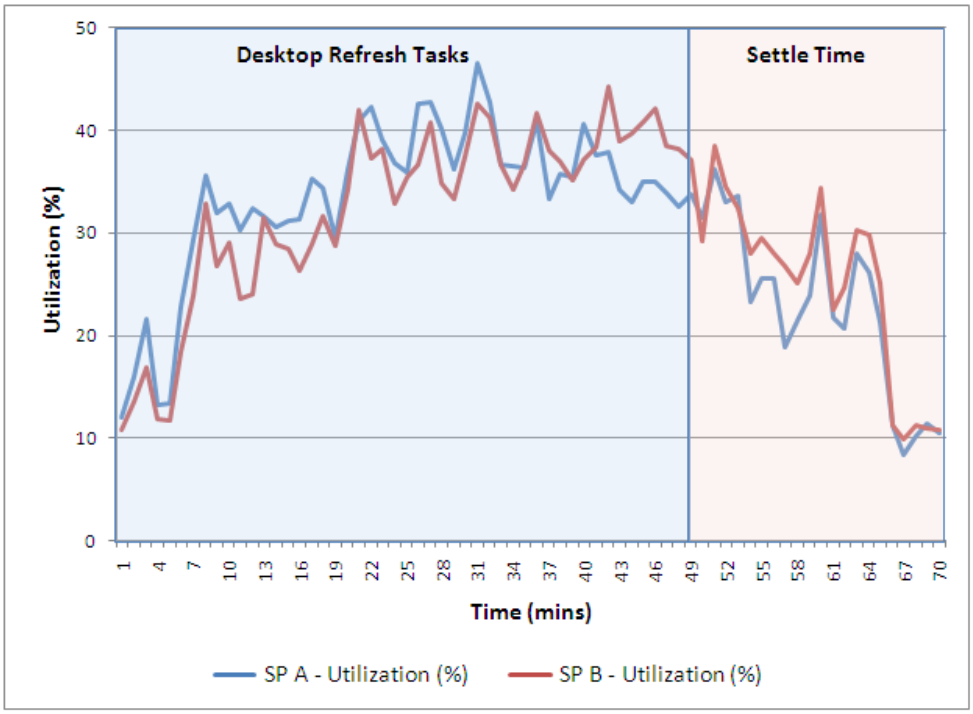

The CPU utilization of the SPs stayed well below 50 percent during the refresh operation. This indicates that the EMC VNX series provides lots of processing headroom.

FAST Cache IOPS

The following graph shows the IOPS serviced from FAST Cache.

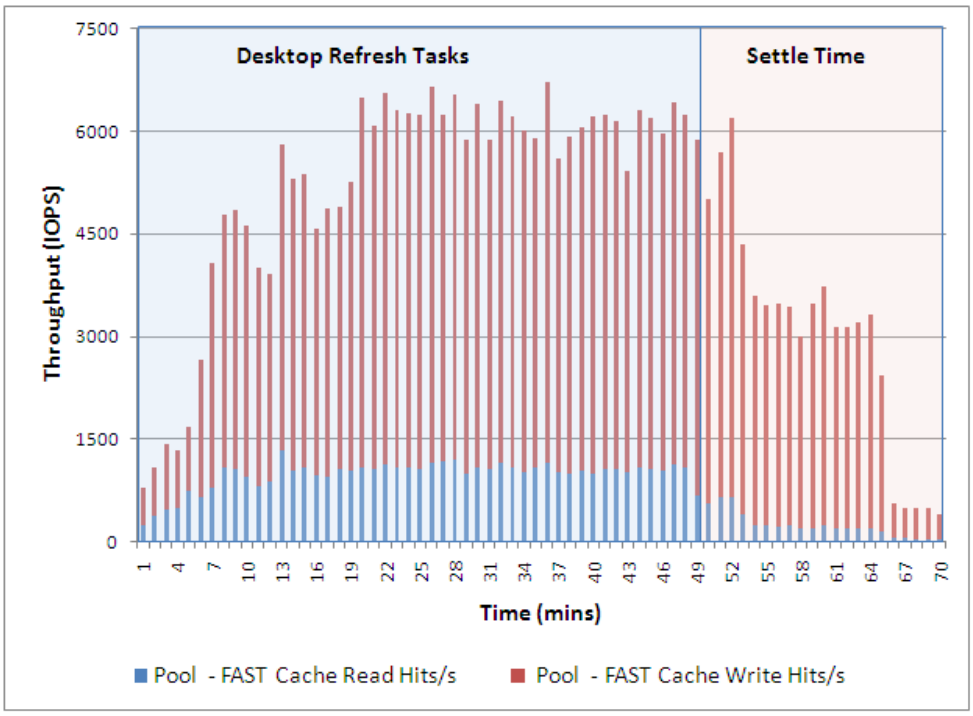

At peak load, FAST Cache serviced over 6,700 IOPS from the linked clone data stores, which is the equivalent of 34 FC disks servicing 200 IOPS each.

ESX CPU load

The following graph shows the CPU load from the ESX servers in the View cluster. A single server is reported because all servers had similar results.

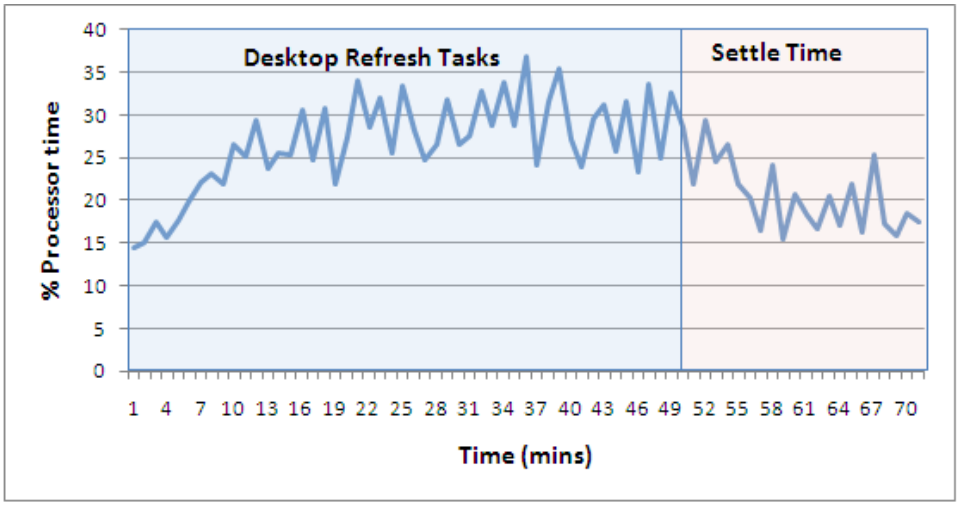

The CPU load of the ESX server was well within the acceptable limits during this test.

ESX disk response time

The following graph shows the Average Guest Millisecond/Command counter, which is shown as GAVG in esxtop. This counter represents the response time for I/O operations issued to the storage array. For the replica, the average of both LUNs hosting the replica storage is shown in the graph.

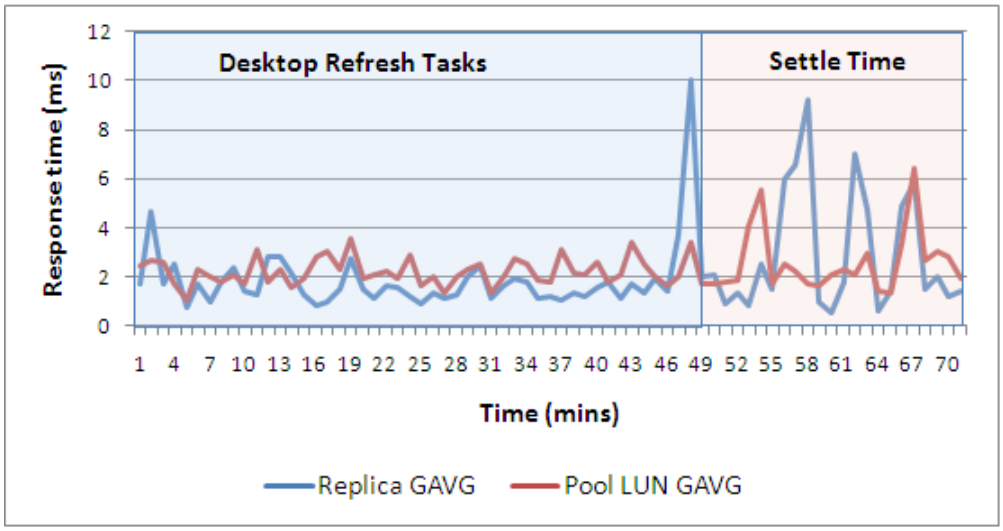

The GAVG values for the EFD replica storage and the linked clone storage on the Pool1_x data stores were well below 10 ms during the peak load. This indicates very good performance under this load.

View recompose results

Test methodology

This test was conducted by creating new pools from the View Manager console. No users were logged in after deploying new desktops.

Overlays are added to the graphs to show when the last power-on task completed and when the IOPS to the pool LUNs achieved a steady state.

A recompose operation deletes existing desktops and creates new desktops. To enhance the readability of the charts and to show the array behavior during high I/O periods, only those tasks involved in creating new desktops were performed and shown in the graphs. Both desktop pools were created simultaneously and took approximately 128 minutes to complete the entire process.

The timeline for the test was as follows:

- 3 to 109 minutes—vCenter Server tasks completed
  - 3 to 10 minutes—Copy the new replica image for pool1
  - 10 to 19 minutes—Create new desktops for pool1
  - 12 to 19 minutes—Copy the new replica image for pool2
  - 19 to 109 minutes—Create new desktops for pool1 and pool2
- 109 to 128 minutes—Settling time for both the pools

In all the graphs, the first highlighted spike of the I/O is the replica copy operation of pool1.

EFD replica disk load

The following graph shows the IOPS from one of the EFDs that contains the replica data store.

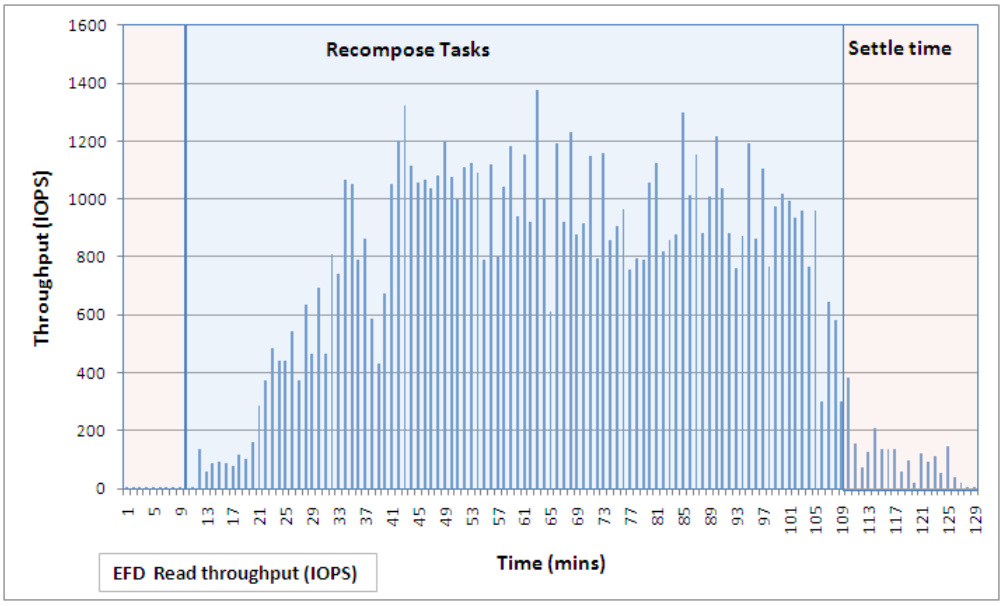

Copying the new replica images caused heavy sequential write workloads on the EFDs. Each EFD serviced nearly 1,400 IOPS at peak load. This indicates that the disks were not loaded heavily during this test.

EFD replica LUN load

The following graph shows the IOPS from the replica LUNs.

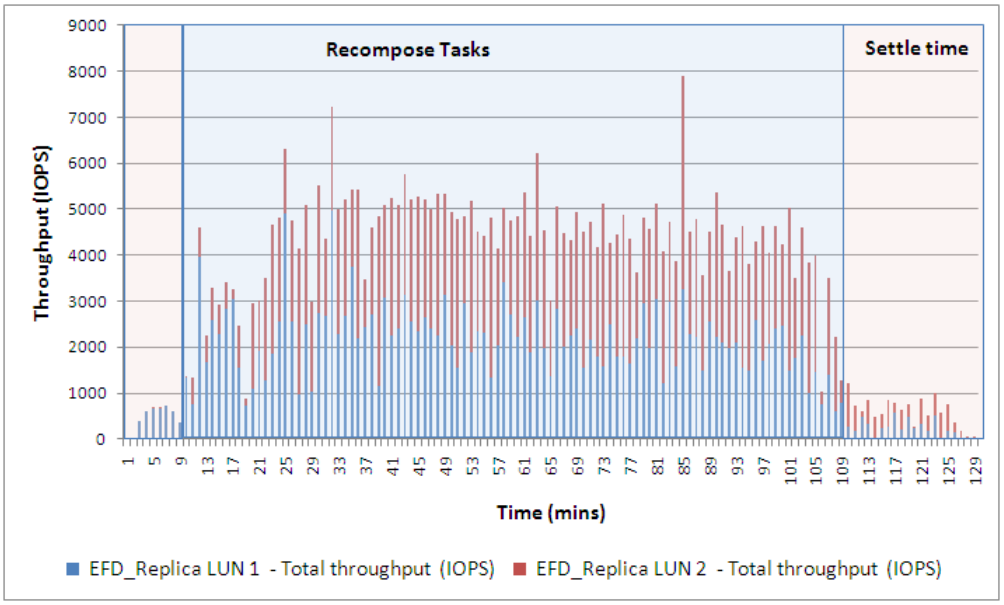

The replica LUNs serviced nearly 7,900 IOPS during peak load.

EFD replica LUN response time

The following graph shows the response time metrics from the replica LUNs.

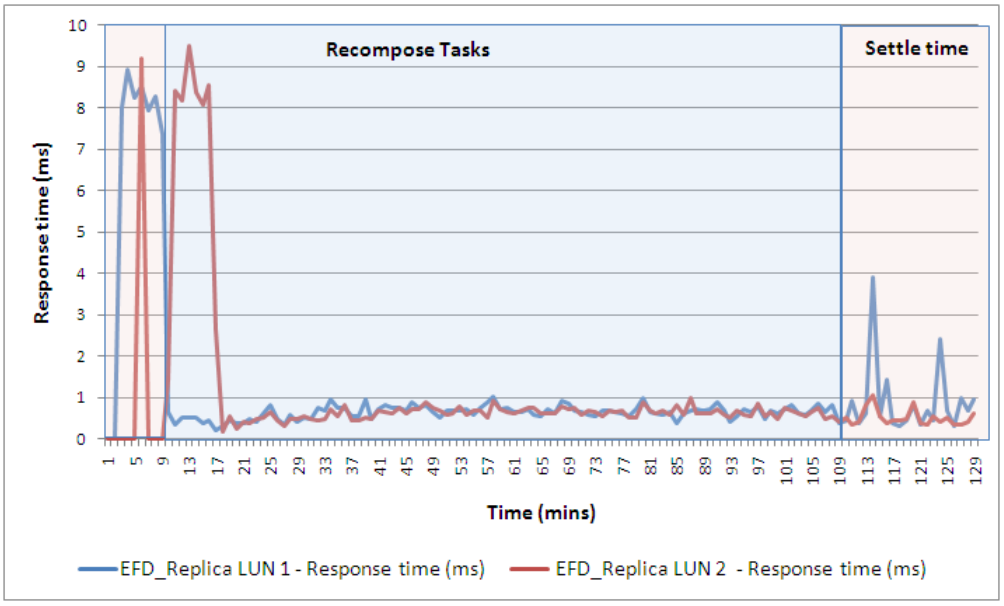

Copying the new replica images caused heavy sequential write workloads on the EFDs, causing small spikes in the response time, that is, up to 10 ms for the EFD LUNs.

Pool individual disk load

The following graph shows the disk IOPS for a single SAS drive in the storage pool that stores the four Pool1_x data stores. Because the statistics from all drives in the pools were similar, only a single drive is reported for clarity and readability of the graph.

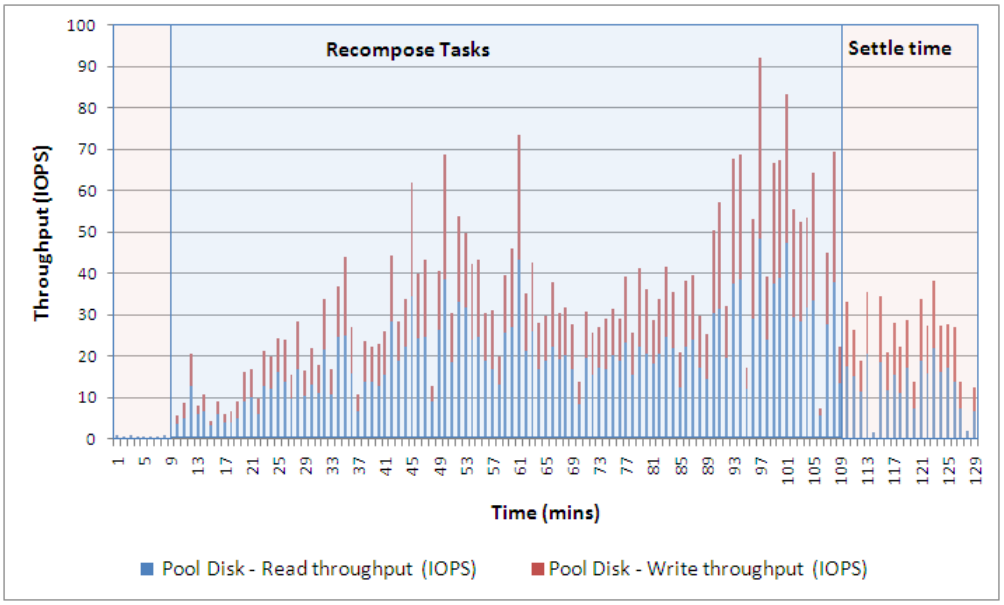

Each drive serviced fewer than 100 IOPS at peak load. Most of the workload was serviced by FAST Cache.

Pool LUN load

The following graph shows the LUN IOPS and response time from the Pool1_3 data store. Because the statistics from all pools were similar, only a single pool is reported for clarity and readability of the graph.

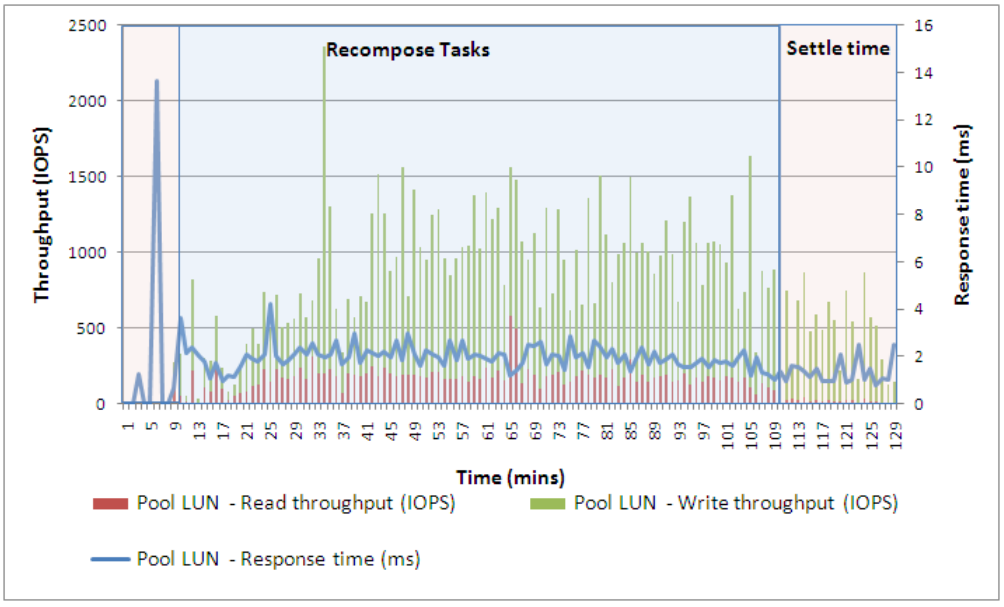

During the test, the LUN response time almost remained within 3 ms during the

creation of new desktops and the data store serviced 2,356 IOPS at peak load.

Storage processor IOPS

The following graph shows the total IOPS served by the storage processor during the test.

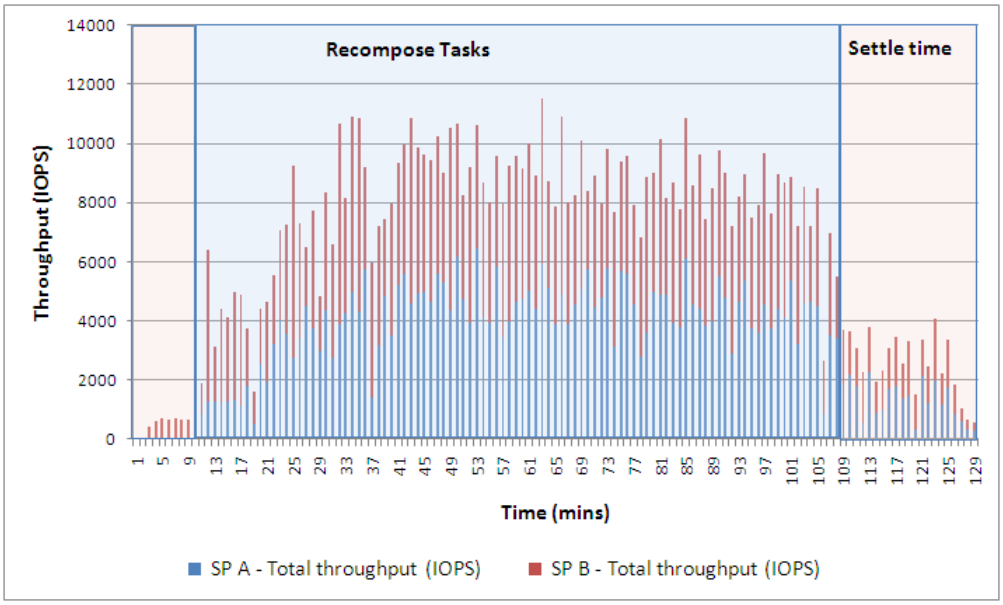

Storage processor utilization

The following graph shows the storage processor utilization during the test.

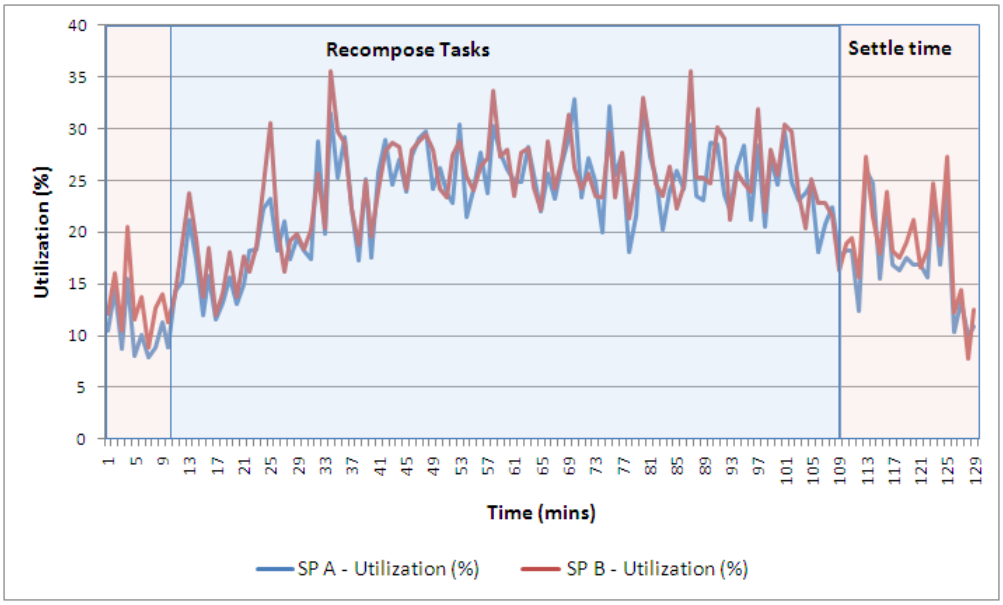

The recompose operation caused moderate CPU utilization during peak load. The VNX series has plenty of scalability headroom for this workload.

FAST Cache I/O The following graph shows the IOPS serviced from FAST Cache during the recompose test.

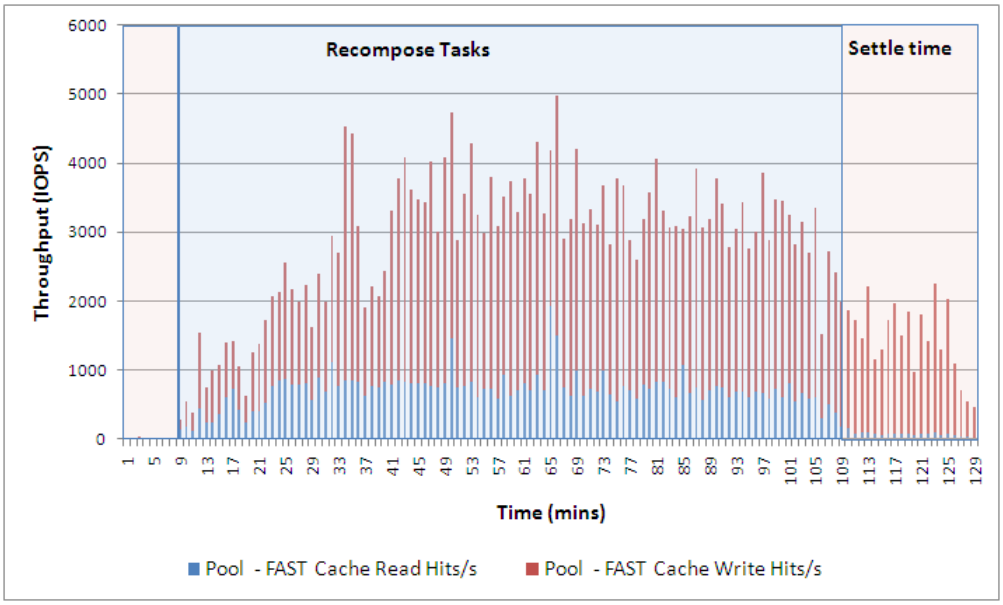

At peak load, FAST Cache serviced nearly 5,000 IOPS from the linked clone data stores, which is the equivalent of 25 FC disks servicing 200 IOPS each.

ESX CPU load The following graph shows the CPU load from the ESX servers in the View cluster. A single server is reported because all servers had similar results.

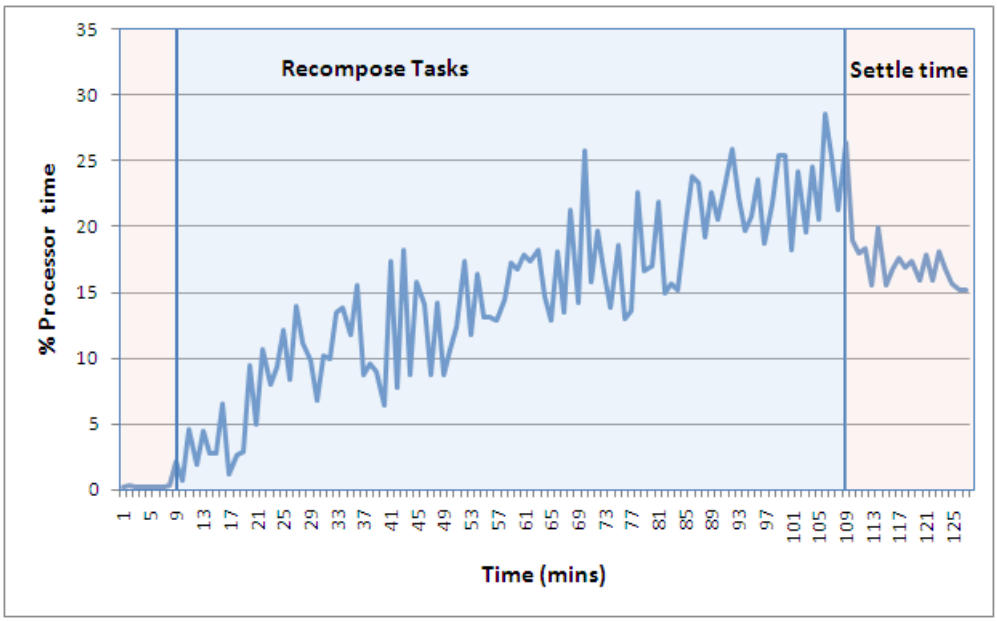

The CPU load on the ESX server was well within acceptable limits during this test.

ESX disk response time

The following graph shows the Average Guest Millisecond/Command counter, which is shown as GAVG in esxtop. This counter represents the response time for I/O operations issued to the storage array. For the replica, the average of both LUNs hosting the replica storage is shown in the graph.

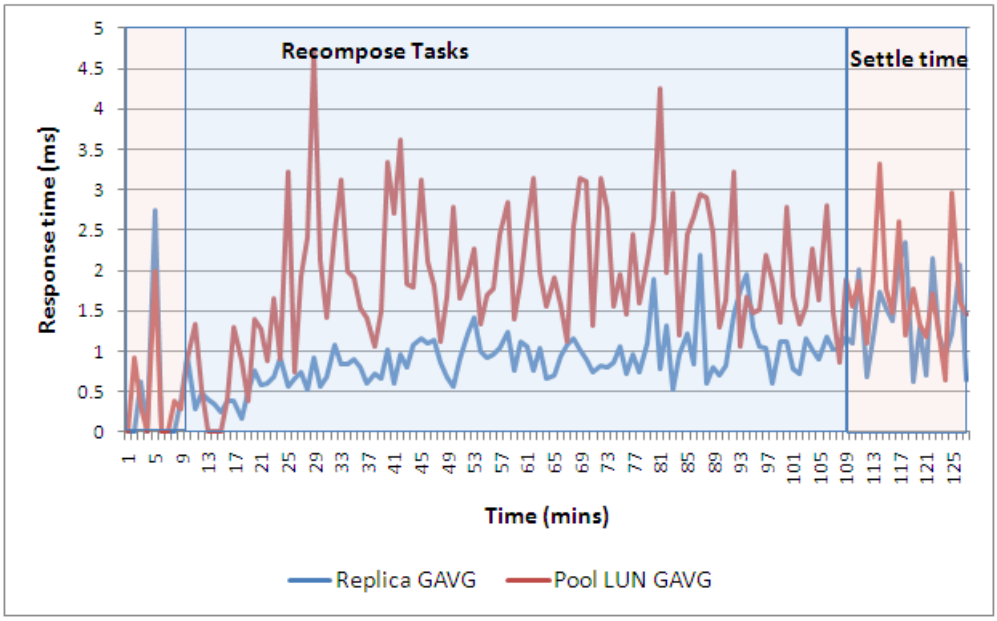

The GAVG values for the EFD replica storage and the linked clone storage on the Pool1_x data stores were below 5 ms. This indicates very good performance under this load.

Antivirus results

Test methodology

This test was conducted by scheduling a full scan of all desktops through a custom script using McAfee 8.7. The desktops were divided into five collections with each collection containing 100 desktops (50 from each pool).The full scans were started over the course of three hours on all desktops with every collection scanning 100 desktops for roughly half an hour.

EFD replica disk load

The following graph shows the IOPS from one of the EFDs that contains the replica data store.

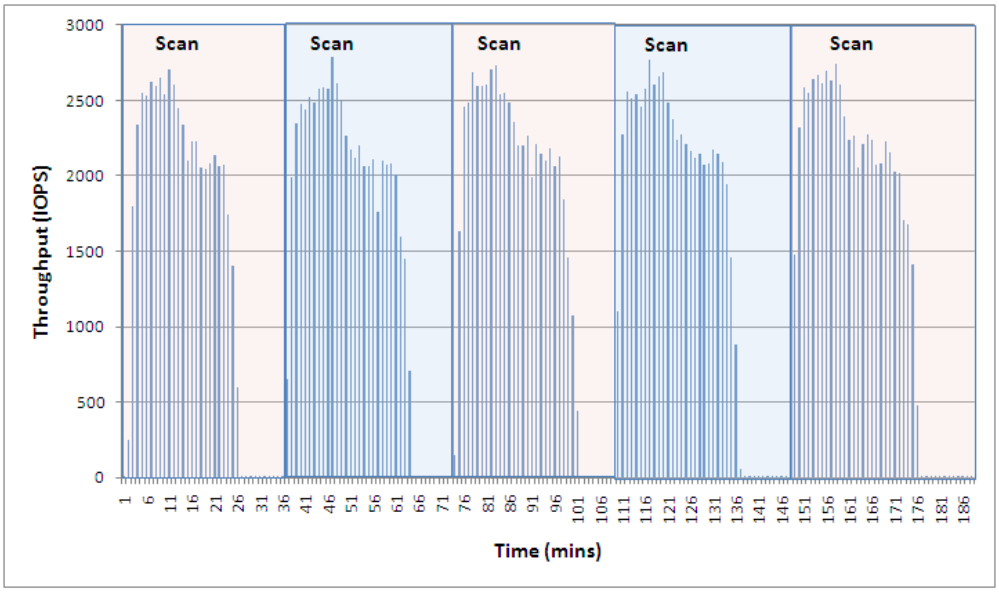

Each EFD serviced nearly 2,800 IOPS at peak load for all five collections.

EFD replica LUN load

The following graph shows the IOPS from the replica LUNs.

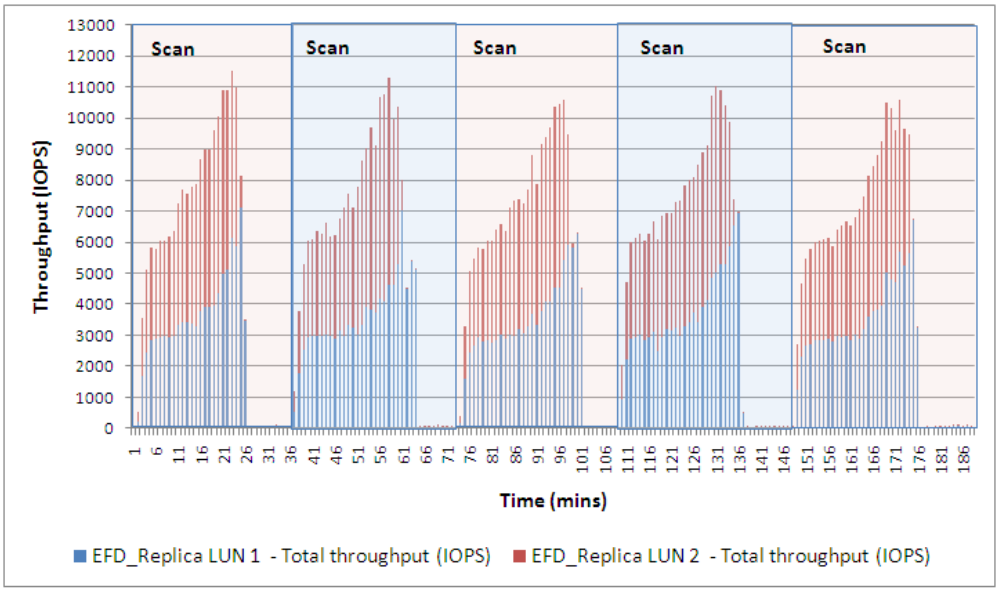

Both replica LUNs serviced a total of 11,500 IOPS during peak load. The load is shared between two replica LUNs during the scan of each collection.

EFD replica LUN response time

The following graph shows the response time metrics from the replica LUNs.

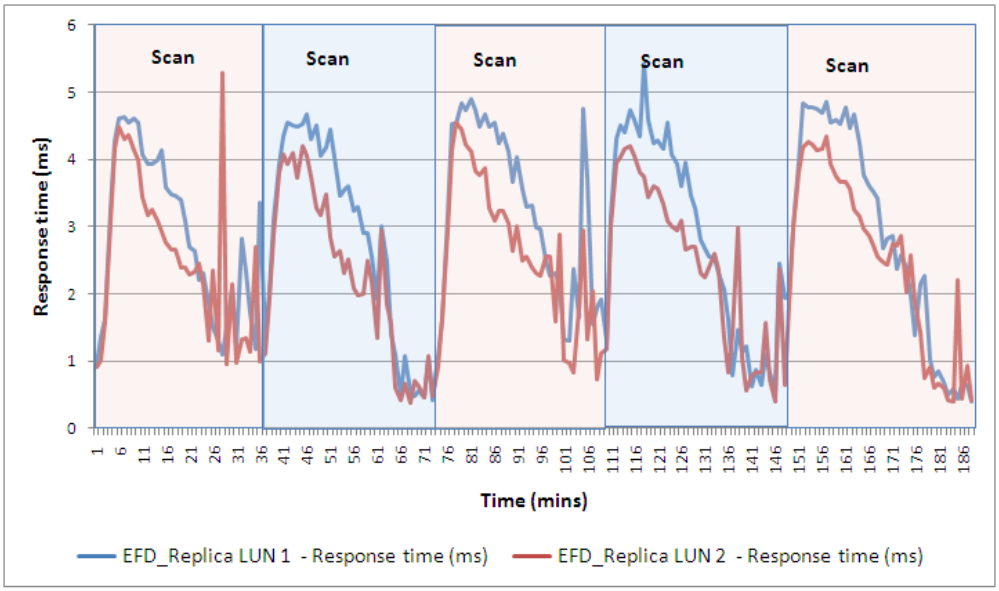

The McAfee scan caused the response time of the replica LUN to spike to 5.4 ms during peak load.

Pool individual disk load

The following graph shows the disk I/O for a single SAS drive in the storage pool that stores the four Pool1_x data stores. Because the statistics from all drives in the pool were similar, only a single drive is reported for clarity and readability of the graph.

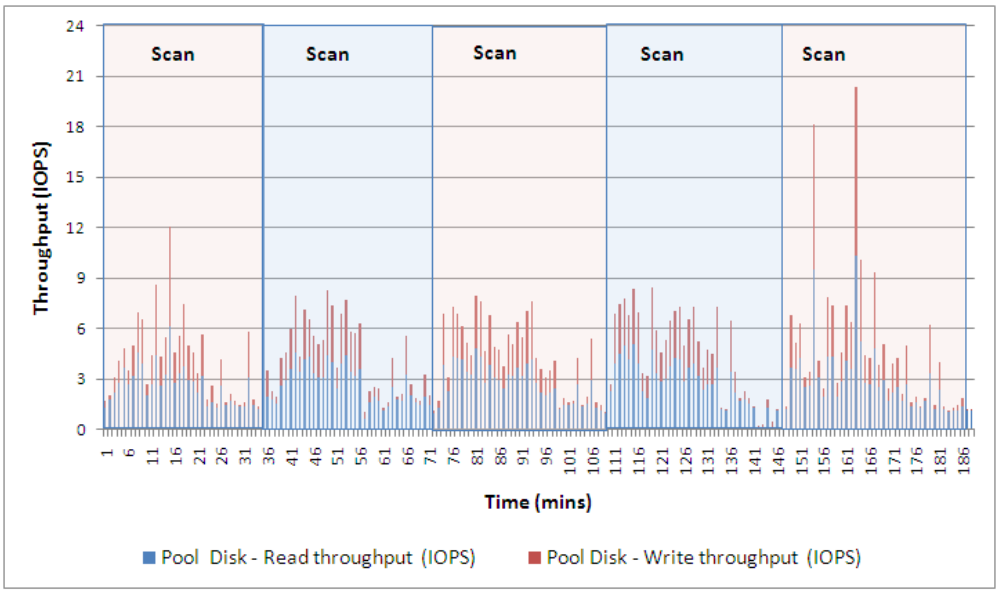

The IOPS serviced by the individual drives in the pool were extremely low because almost all I/O requests were serviced either by FAST Cache or EFD replicas.

Pool LUN load

The following graph shows the LUN IOPS and response time from the Pool1_3 data store. Because the statistics from all the pools were similar, only a single pool is reported for clarity and readability of the graph.

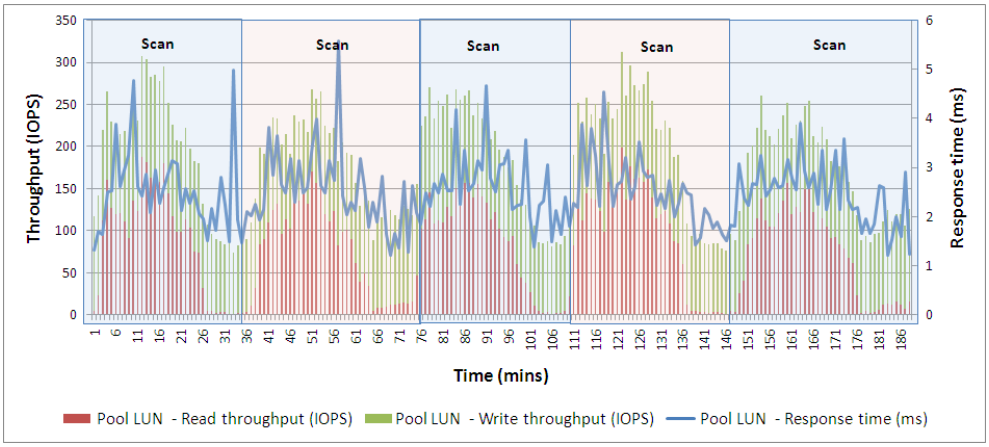

During peak load, the LUN response time remained within 6 ms and the data store serviced over 300 IOPS. The majority of the read I/O was served by the replica LUNs and not by the pool LUN.

Storage processor IOPS

The following graph shows the total IOPS serviced by the storage processor during the test.

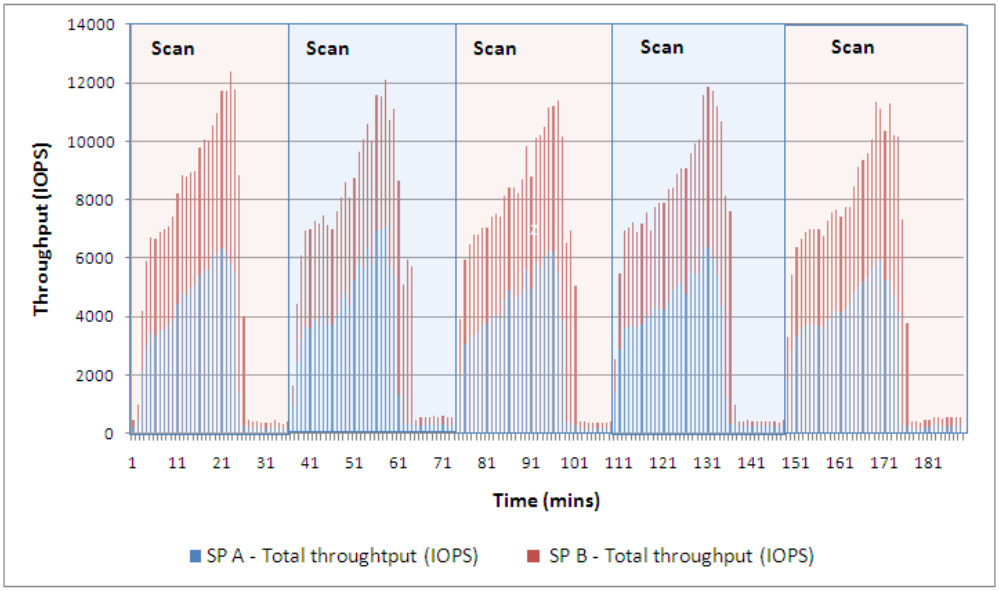

Storage processor utilization

The following graph shows the storage processor utilization during the test.

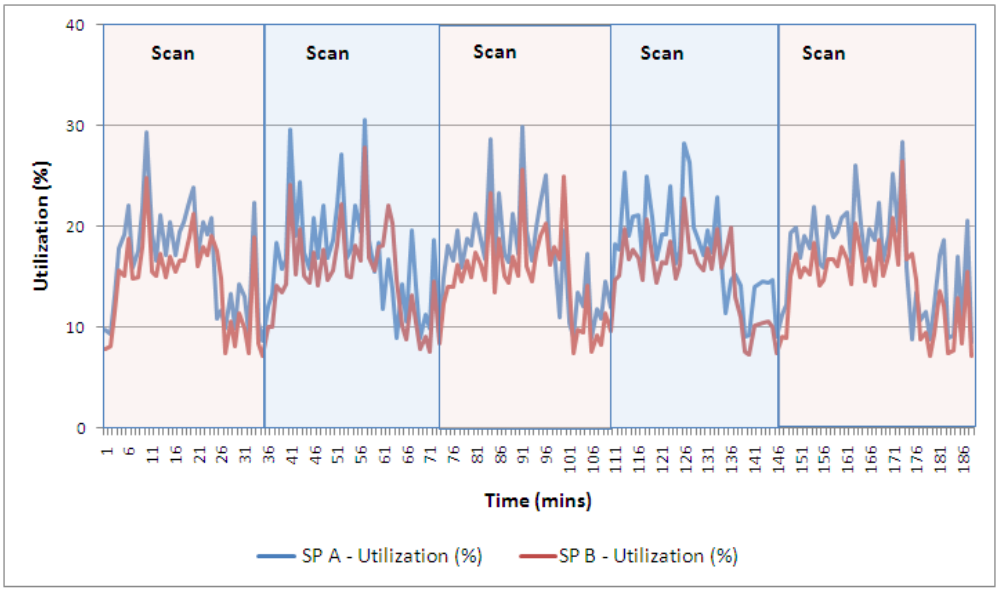

The antivirus scan operations caused moderate CPU utilization during peak load. The load is shared between two SPs during the scan of each collection. The EMC VNX series has plenty of scalability headroom for this workload.

FAST Cache IOPS

The following graph shows the IOPS serviced from FAST Cache during the test.

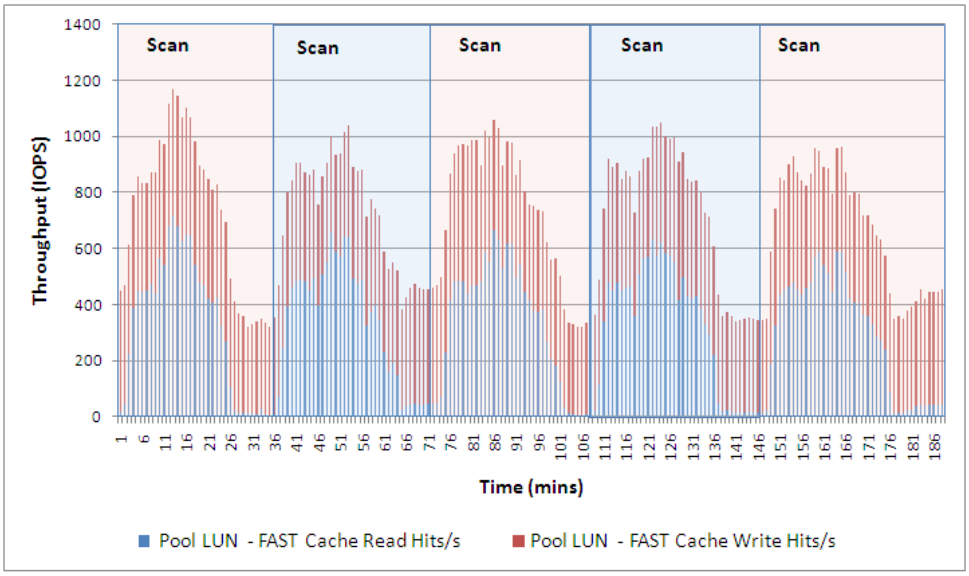

At peak load, the FAST Cache serviced over 1,100 IOPS from the linked clone data stores, which is the equivalent of six FC disks serving 200 IOPS each. The majority of the read operations are serviced from the replica during this test.

ESX CPU load

The following graph shows the CPU load from the ESX servers in the View cluster. A single server is reported because all servers had similar results.

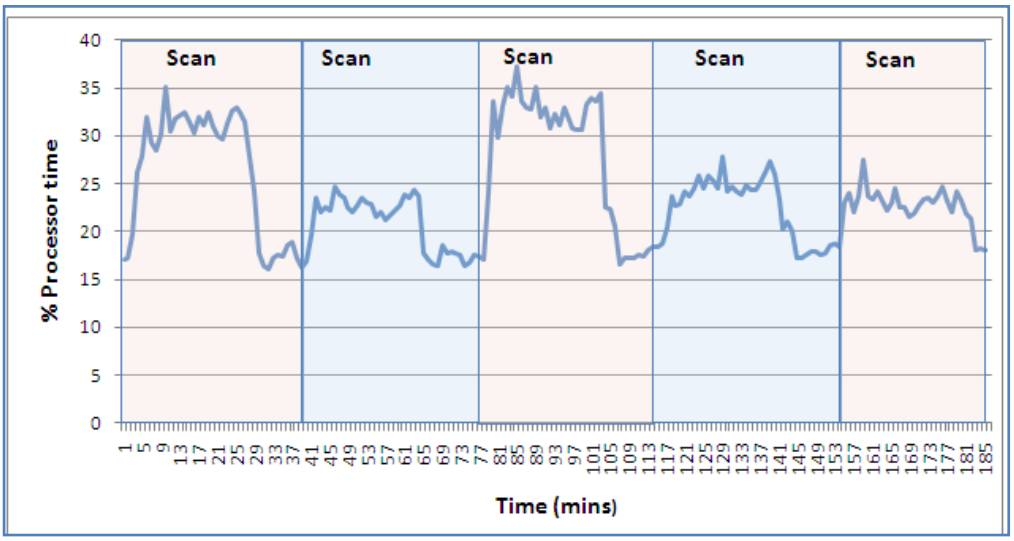

The CPU load on the ESX server was well within acceptable limits during this test.

ESX disk response time

The following graph shows the Average Guest Millisecond/Command counter, which is shown as GAVG in esxtop. This counter represents the response time for I/O operations issued to the storage array. For the replica, the average of both LUNs hosting the replica storage is shown in the graph.

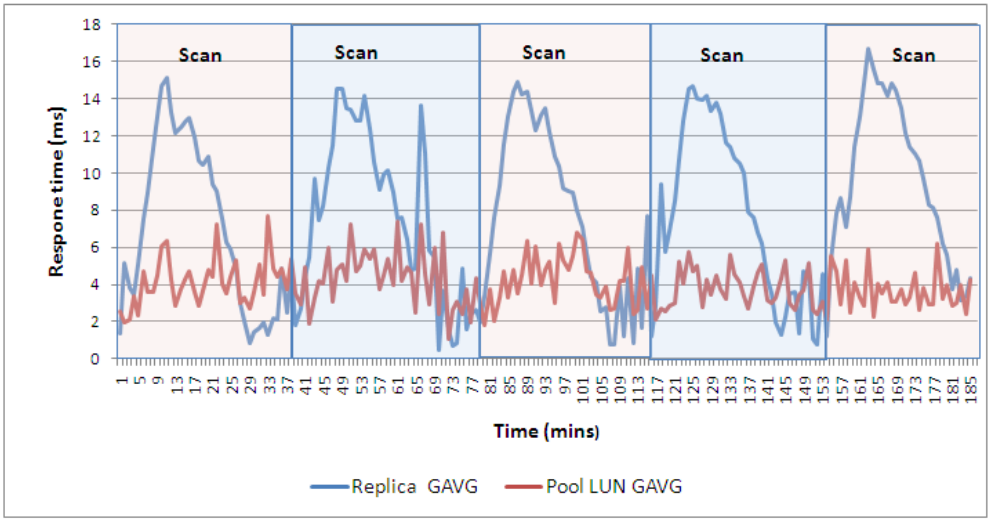

The peak GAVG value for the EFD replica storage was well below 20 ms. The replicas performed an enormous amount of read I/O operations during this test. The GAVG values for the linked clone storage on the Pool1_x data stores were below 8 ms.

Patch install results

Test methodology

This test was performed by pushing five security updates to all desktops using Microsoft System Center Configuration Manager (SCCM). The desktops were divided into five collections with each containing 100 desktops (50 from each pool).

The collections were configured to install updates in a 15-minute staggered schedule. This caused all patches to be installed in nearly 1 hour and 15 minutes.

EFD replica disk load

The following graph shows the IOPS from one of the EFDs that contains the replica data store.

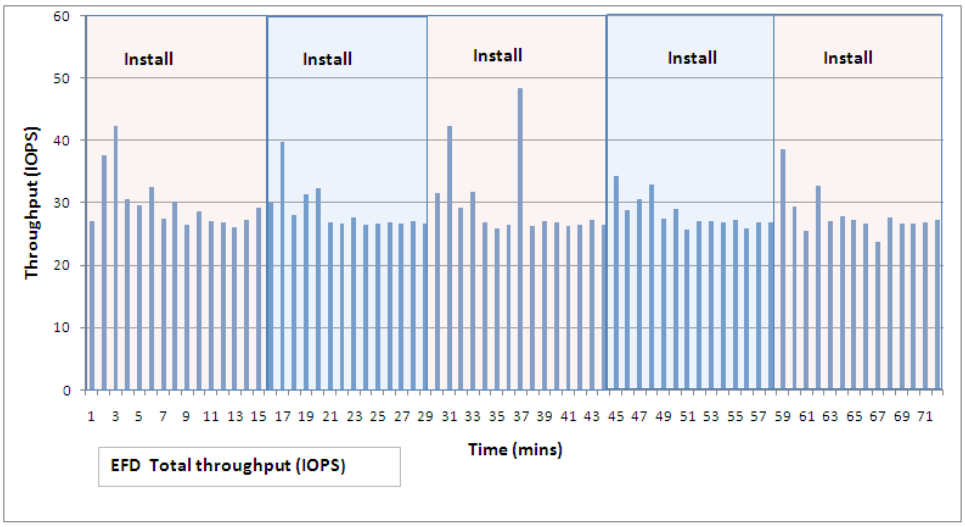

The load on the EFDs was minimal. At peak load, each EFD serviced less than 50 IOPS.

EFD replica LUN load

The following graph shows the IOPS from the replica LUNs.

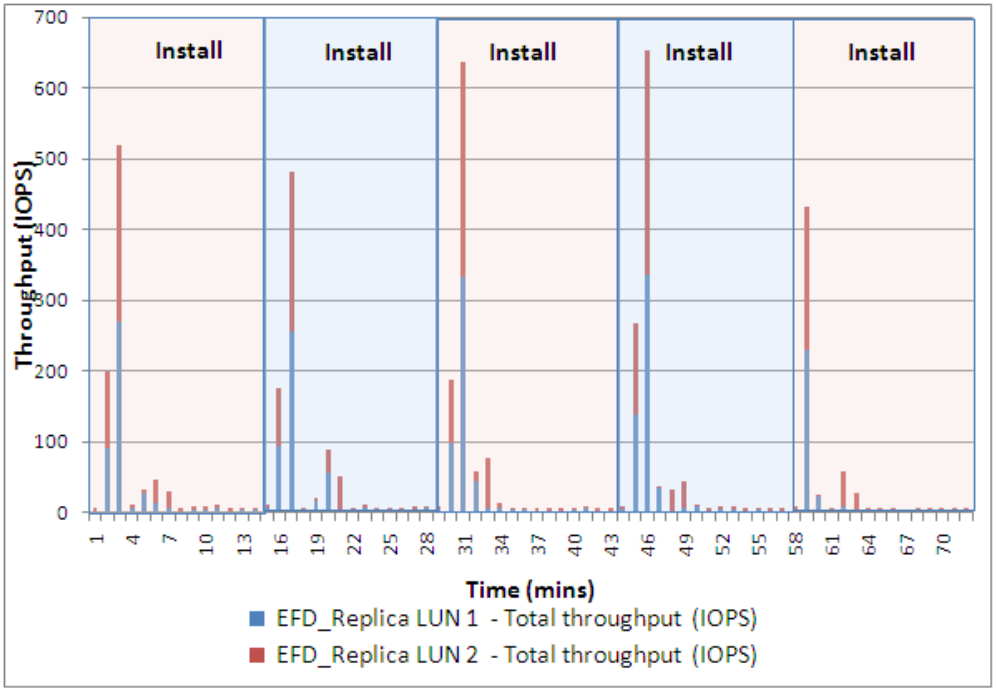

The replica LUNs serviced over 650 IOPS during peak load.

EFD replica LUN response time

The following graph shows the response time metrics from the replica LUNs.

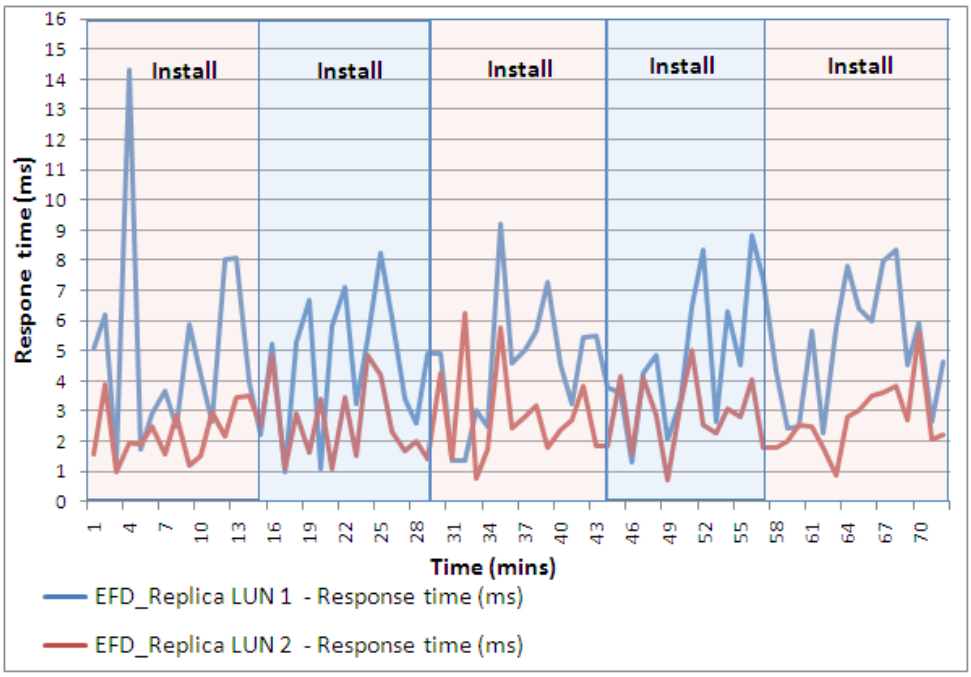

The response time values of the replica LUNs were below 10 ms, apart from the initial peak that crossed 14 ms.

Pool individual disk load

The following graph shows the disk IOPS for a single SAS drive that consists of the storage pool that stores the four Pool1_x data stores. Because the statistics from all drives in the pool are similar, the statistics of a single drive are shown in the graphs for clarity and readability.

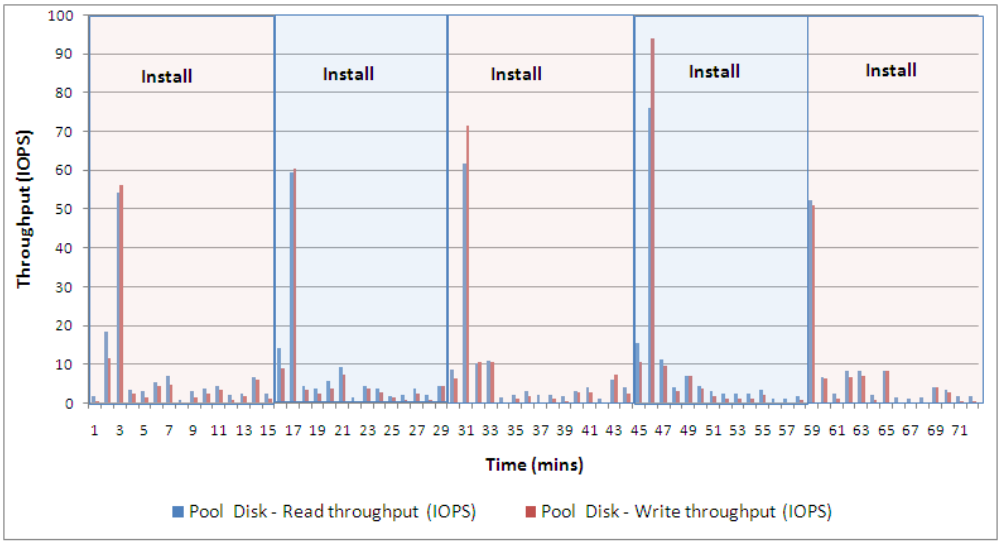

The drives did not get saturated because the majority of IOPS was serviced by FAST Cache.

Pool LUN load

The following graph shows the LUN IOPS and response time from the Pool1_3 data

store. Because the statistics from all drives in the pool were similar, the statistics of a single drive are shown in the graphs for clarity and readability.

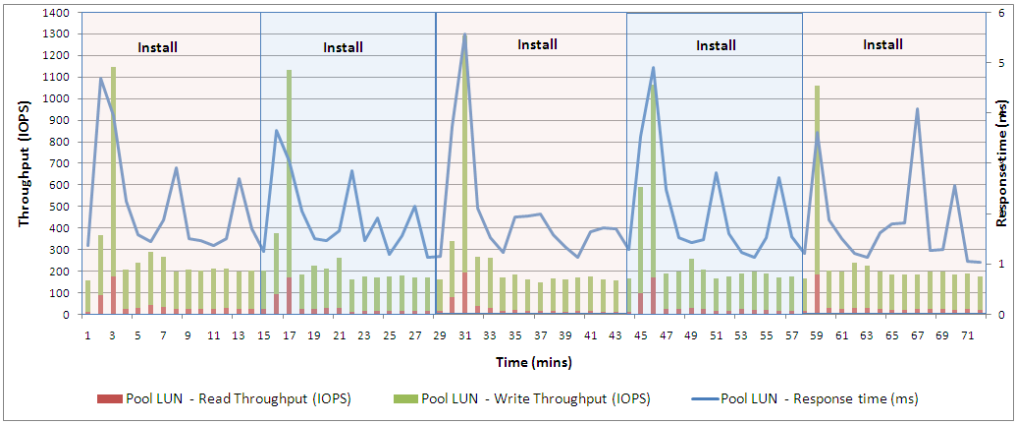

During peak load, the LUN response time was below 6 ms and the data store serviced nearly 1,300 IOPS during peak load.

Storage processor IOPS

The following graph shows the total IOPS serviced by the storage processor during the test.

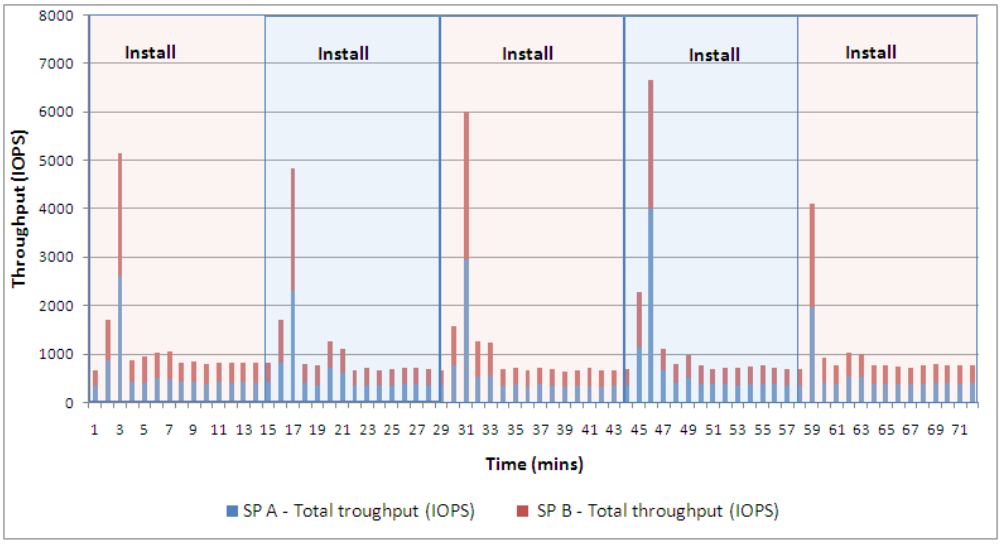

During peak load, the storage processors serviced over 6,000 IOPS. The load is shared between two SPs during the patch install operation of each collection.

Storage processor utilization

The following graph shows the storage processor utilization during the test.

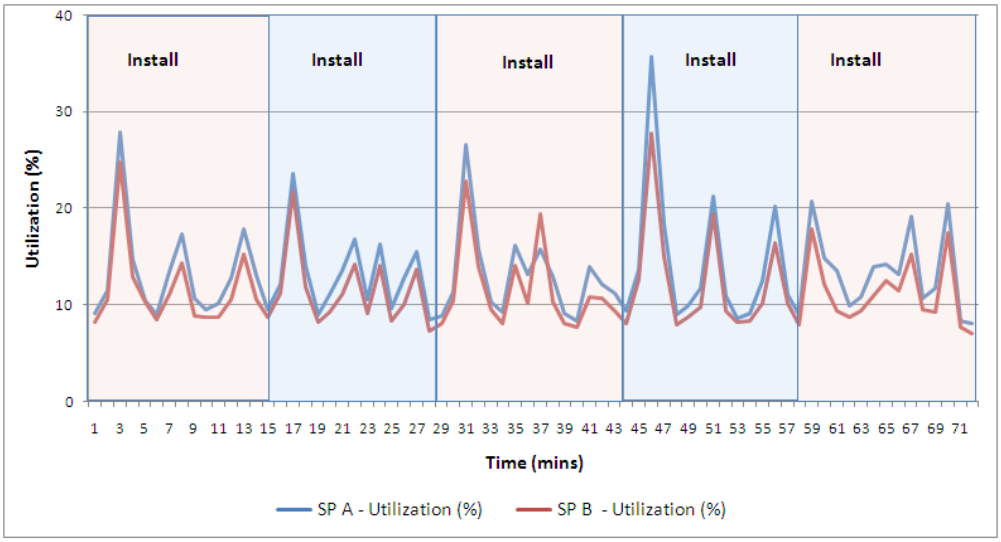

The patch install operations caused moderate CPU utilization during peak load. The EMC VNX series has plenty of scalability headroom for this workload.

FAST Cache IOPS

The following graph shows the IOPS serviced from FAST Cache during the test.

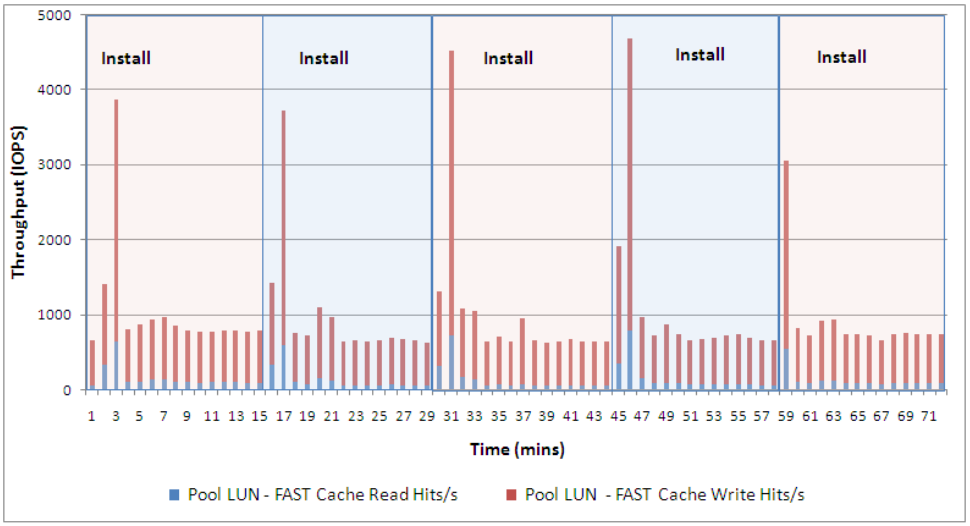

FAST Cache serviced over 4,500 IOPS at peak load from the linked clone data stores, which is the equivalent to 23 FC disks servicing 200 IOPS each.

ESX CPU load

The following graph shows the CPU load from the ESX servers in the View cluster. Because all servers had similar results, the results from a single server are shown in the graph.

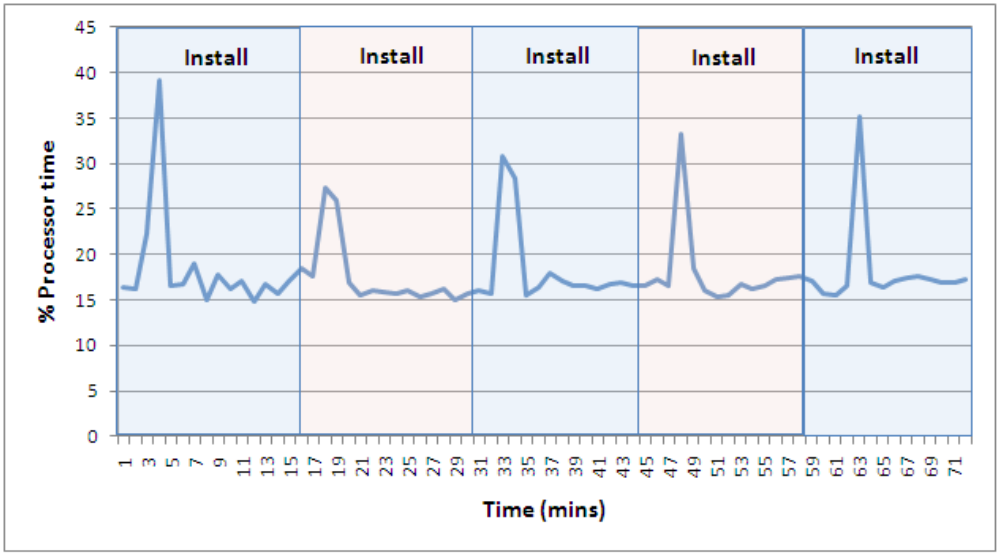

The ESX server CPU load was well within the acceptable limits during the test.

ESX disk response time

The following graph shows the Average Guest Millisecond/Command counter, which is shown as GAVG in esxtop. This counter represents the response time for the I/O operations issued to the storage array. The average of both LUNs hosting the replica storage is shown in the graph.

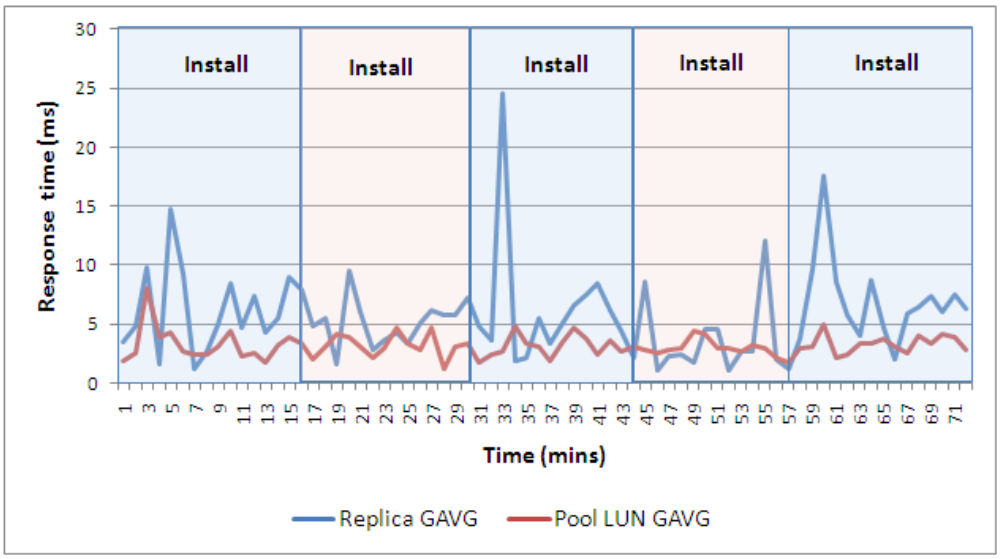

The GAVG values for linked clone storage on the Pool1_x data stores were well below 5 ms for almost all data points. The GAVG value for the replica data store peaked to 25 ms.

Login VSI results

Test methodology

This test was conducted by scheduling 500 users to connect over a Remote Desktop Protocol (RDP) connection in a 35-minute window and start the Login VSI-medium workload. This workload was run for two hours in a steady state to observe the load on the system.

EFD replica disk load

The following graph shows the IOPS from one of the EFDs that contains the replica data store.

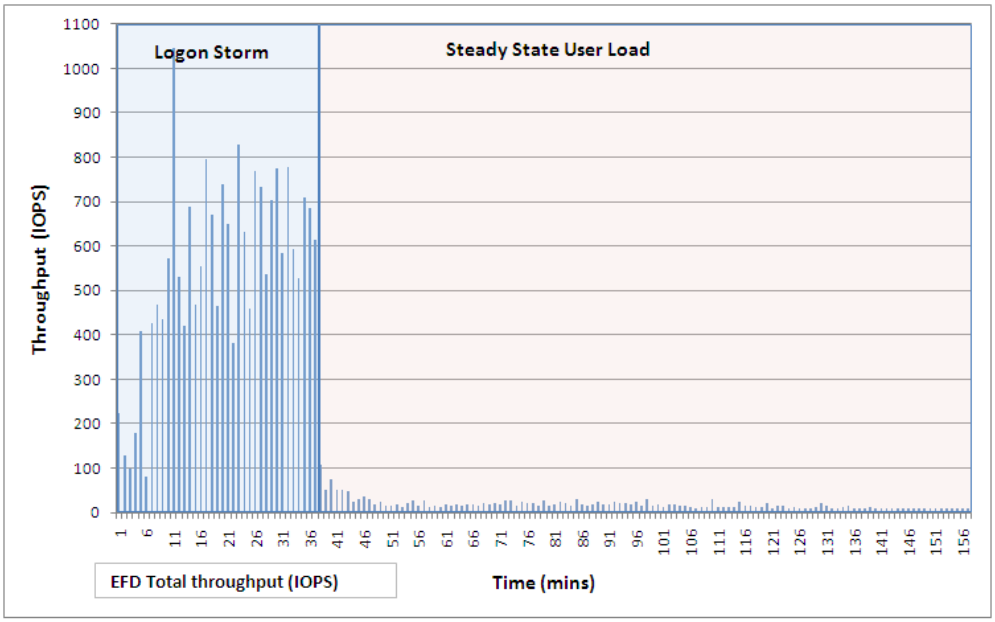

During logon to all desktops, the peak read load on the EFD reached over 1,000 IOPS.

EFD replica LUN load

The following graph shows the IOPS from the replica LUNs.

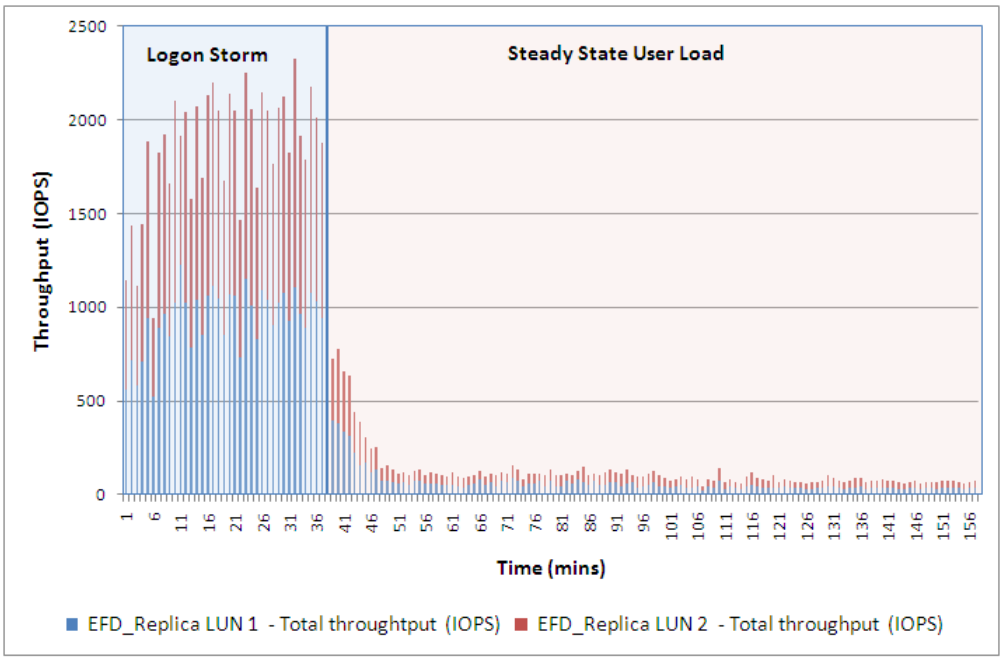
Logon Storm
Steady State User Load
Throughput (IOPS)
Time (mins)
EFD_Replica LUN 1 - Total throughput (IOPS)
EFD_Replica LUN 2 - Total throughput (IOPS)

The replica LUNs serviced over 2,300 IOPS during peak load.

EFD replica LUN response time

The following graph shows the response time metrics from the replica LUNs.

.

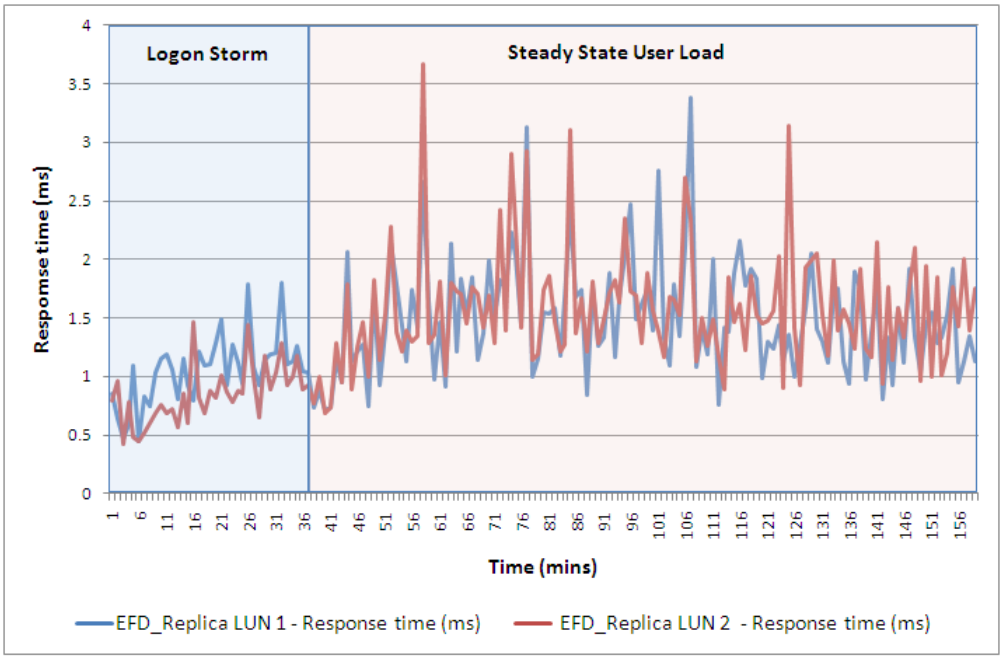
Logon Storm
Steady State User Load
Respomse time (ms)
Time (mins)
EFD_Replica LUN 1 - Response time (ms)
EFD_Replica LUN 2 - Response time (ms)

During steady state load, the response time was almost below 3.5 ms.

Pool individual disk load

The following graph shows the disk IOPS from the Pool1_3 data store. Because the statistics from all the pools were similar, only a single pool is reported for clarity and readability of the graphs.

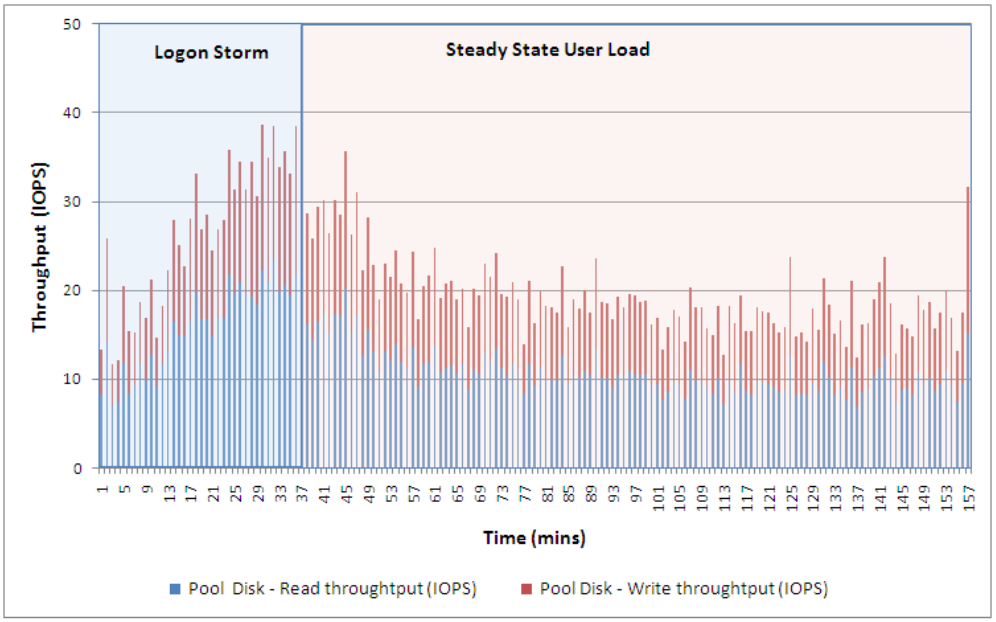

During peak load, the SAS disk serviced less than 40 IOPS.

Pool LUN load

The following graph shows the LUN IOPS and response time from the Pool1_3 data store. Because the statistics from all pools were similar, only a single pool is reported for clarity and readability of the graphs.

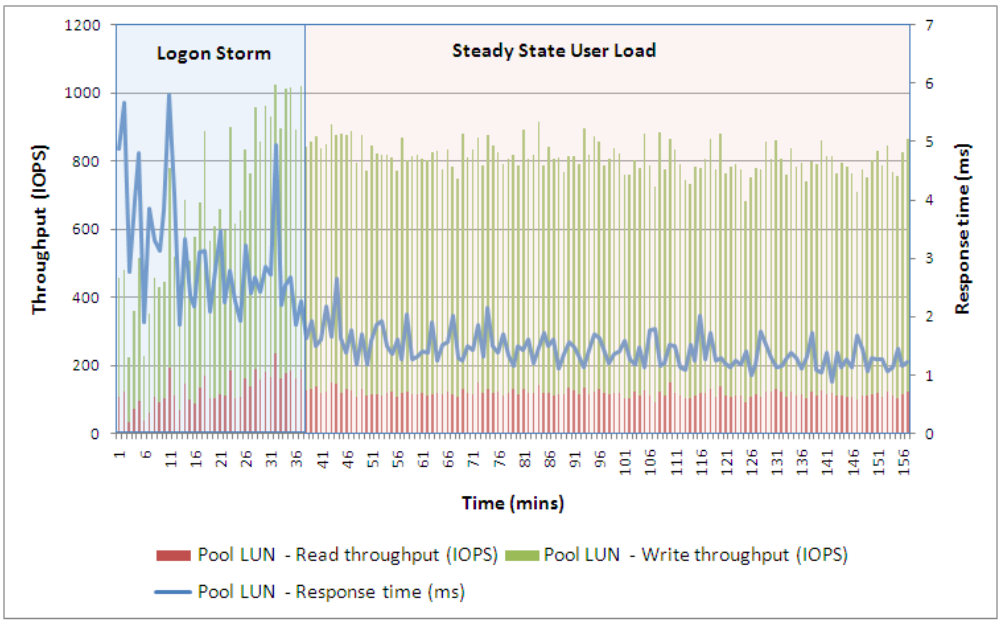

During peak load, the LUN response time remained within 6 ms and the data store serviced over 1,000 IOPS.

Storage processor IOPS

The following graph shows the total IOPS serviced by the storage processor during the test.

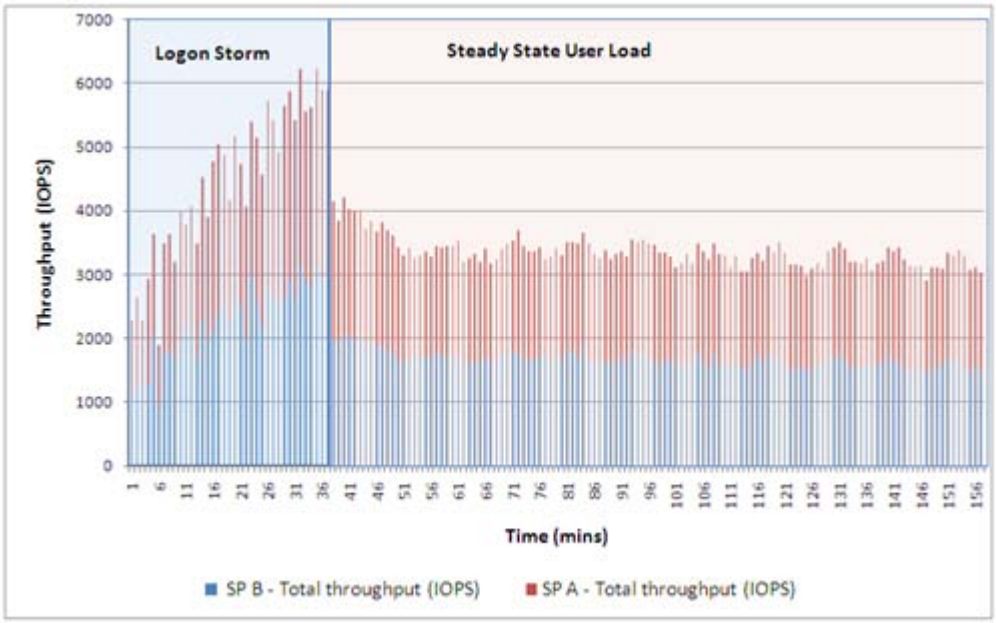

Storage processor utilization

The following graph shows the storage processor utilization during the test.

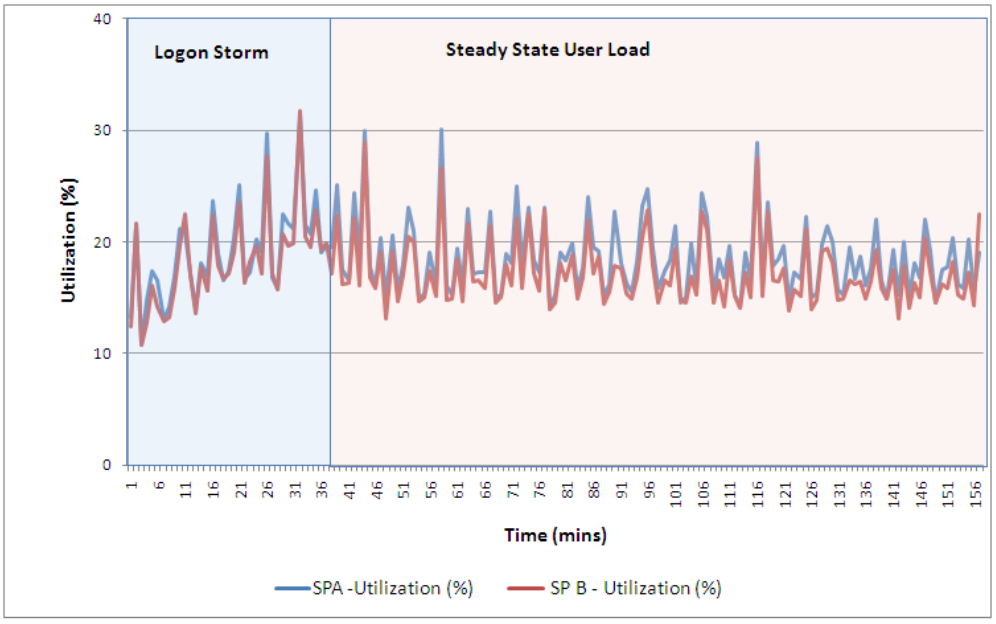

The storage processor utilization peaks at a little over 30 percent during the logon storm. The load is shared between two SPs during the VSI load test.

FAST Cache IOPS

The following graph shows the IOPS serviced from FAST Cache during the test.

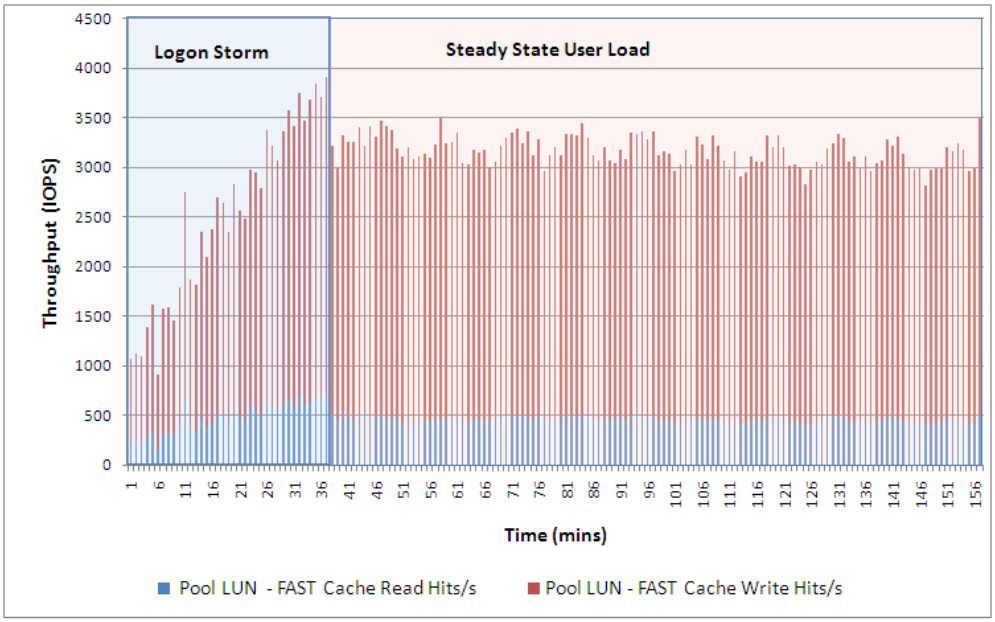
Logon Storm
Steady State User Load
Throughput (IOPS)
Time (mins)
Pool LUN - FAST Cache Read Hits/s
Pool LUN - FAST Cache Write Hits/s

At peak load, FAST Cache serviced nearly 4,000 IOPS from the linked clone data stores, which is the equivalent of 20 FC disks serving 200 IOPS each.

ESX CPU load

The following graph shows the CPU load from the ESX servers in the View cluster. A single server is reported because all servers had similar results.

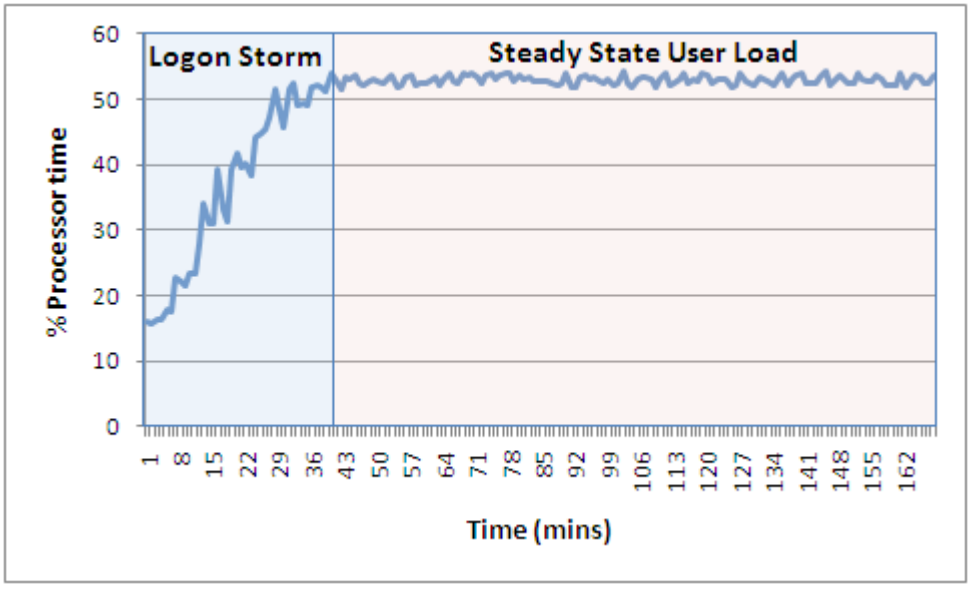

The CPU load on the ESX server was well within the acceptable limits during the test.

ESX disk response time

The following graph shows the Average Guest Millisecond/Command counter, which is shown as GAVG in esxtop. This counter represents the response time for I/O operations issued to the storage array. The average of both LUNs hosting the replica storage is shown in the graph.

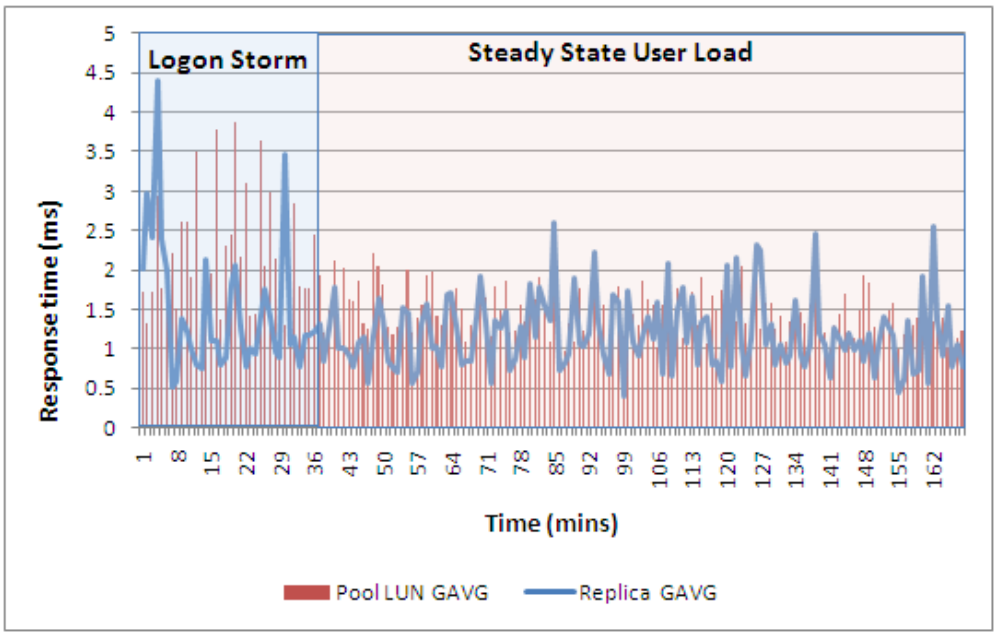

The GAVG values for the EFD replica storage and the linked clone storage on the Pool1_x data stores were well below 4.5 ms during the peak load.